The impact of demand parameter uncertainty on the bullwhip effect

Erica Pastore[a,*], Arianna Alfieri[a], Giulio Zotteri[a], John E. Boylan[b]

[a]*Politecnico di Torino, Torino, Italy*
[b]*Centre for Marketing Analytics and Forecasting, Lancaster University, Lancaster, UK*

Abstract

The bullwhip effect is a very important issue for supply chains, impacting on costs and effectiveness. Academic researchers have studied this phenomenon and modelled it analytically, showing that it affects many real world industries. The analytical models generally assume that the final demand process and its parameters are known. This paper studies a two-echelon single-product supply chain with final demand distributed according to a known AR(1) process but with unknown parameters. The results show that the bullwhip effect is affected by unknown parameters and is influenced by the frequency with which parameter estimates are updated. For unknown parameters, the strength of the bullwhip effect is also influenced by the number of demand observations available to estimate the parameters. Furthermore, a negative autoregressive parameter does not always imply an anti-bullwhip effect when the parameters are unknown. An analytical approximation is proposed to mitigate the poor accuracy of existing models when the parameters of an AR(1) process are unknown, forecasts are updated but parameter estimates remain unchanged.

Keywords: Inventory, Supply Chain Management, Bullwhip Effect, Demand Variability

1. Introduction

Supply chain effectiveness and costs are usually affected by demand variability, especially in the upstream echelons. It has been shown that demand variability tends to be amplified moving upstream in the supply chain (the so called bullwhip effect phenomenon), and that this amplification tends to increase supply chain costs and to damage the service to the final customers (Lee et al., 1997). Higher demand variability implies higher uncertainty in inventory processes and, as a consequence, higher inventory costs.

*Corresponding author
Email address: erica.pastore@polito.it (Erica Pastore)

Many industrial sectors have been found to be affected by the bullwhip effect: the American automobile sector (Blanchard, 1983), the machine tool industry (Anderson et al., 2000), the computer and semi-conductor industry (Terwiesch et al., 2005), to cite a few. Moreover, many cases of supply chains have been studied in the literature to identify the bullwhip and understand its main drivers. Besides the seminal case studies of Procter and Gamble (Lee et al., 1997) and Barilla (Hammond, 1994), demand variability amplification has been found in big retail organisations such as Tesco and Wal-Mart (Disney, 2007; Gill & Abend, 1997). Also, automotive companies suffer from the bullwhip effect when moving from the final customers to the assemblers and the producers (Bray & Mendelson, 2015; Edgehill et al., 1987; Pastore et al., 2019).

To reduce bullwhip related costs, companies may address the issues of identifying the set of products (among their whole assortment) that are affected by the bullwhip (thus the ones to focus on) and of quantifying the size and the nature of the investment to undertake to reduce this effect. Information sharing, incentive design, collaboration and integration are some of the main remedial strategies to reduce variability amplification (Hosoda et al., 2008; Lee & Whang, 2000). Also, by knowing the main drivers influencing such demand variability amplification, companies may want to invest in adjusting these drivers to decrease it. The academic literature presents many contributions on the analysis of the influencing factors of the bullwhip effect (e.g., Chen et al., 2000a; Lee et al., 1997; Rong et al., 2008). These studies exploit analytical methods to investigate the above mentioned factors and usually make some assumptions about the inventory characteristics. Even though the assumptions do not always hold in real contexts, the analytical models may still yield insights for real world companies to understand the drivers of bullwhip related costs.

Although it is useful to identify bullwhip drivers, it is also necessary to quantify their effect. It is here that the assumptions made in analytical models may make them misleading. For example, a model may over-estimate the benefits of information sharing, which could lead to a misguided investment in implementing new information systems. The analytical models in the literature generally assume that the final demand process and its parameters are known (e.g., Chen et al., 2000a,b; Lee et al., 2000). However, in the real world, neither the generating process of final demand nor its parameters are known with certainty. This paper aims to provide more accurate estimates of the bullwhip effect, for AR(1) final demand, by breaking one of the assumptions usually made regarding the final demand process. Specifically, it assumes the parameters of the demand distribution to be unknown. The impact of the uncertainty of the parameters of the final demand process on the bullwhip effect is investigated, with particular attention to its interactions with long lead times and small sets of demand observations to estimate the parameters.

The paper is organised as follows. Section 2 reviews the literature on the bullwhip effect. Section 3 highlights the contributions of the paper. The inventory system and its processes are described in Section 4. Section 5 presents a simulation experiment to check whether not knowing the parameters affects the value of the bullwhip effect. In Section 6, an approximate analytical quantification of the bullwhip effect is proposed for the case of an AR(1) demand process with unknown parameters. The accuracy of the analytical approximation is investigated in Section 7. Section 8 summarises the paper's conclusions, and includes a discussion on limitations and further research opportunities.

2. Literature Review

The phenomenon of demand variability amplification was first investigated in the late 1950s (Forrester, 1958). Since then, many researchers have studied the bullwhip effect both from empirical and analytical standpoints, modelling this phenomenon, understanding its influencing factors and looking for possible remedies. This paper focuses on the relationship between the bullwhip and demand parameter uncertainty, taking into account lead times and target Cylcle Service Levels, and develops an analytical model for the bullwhip. This addresses the issue of demand signal processing, whereby demand variance is amplified as a consequence of the inherent demand properties and inventory rules. Many authors have studied the relationship between the bullwhip and demand signal processing, by developing analytical models to quantify demand variability propagation, by investigating the effect of information sharing, and by modifying the inventory processes to reduce the bullwhip. As this paper addresses the propagation of demand variability, contributions investigating other approaches are not reviewed in the following. The same holds for the contributions related to other causes of the bullwhip such as price fluctuations, order batching and shortage gaming.

Huge efforts have been made to formulate analytical expressions for the bullwhip effect, with the aim of understanding the behaviour of demand variability propagation in supply chains with various characteristics. The lead time has been proved to be a critical factor in influencing the bullwhip effect (Lee et al., 2000).

Usually, when considering the theoretical analyses of the bullwhip, the Order-Up-To (OUT) policy is used to model the planning process of inventory systems (e.g., Boute et al., 2014; Chen et al., 2000a; Disney et al., 2006; Lee et al., 2000; Luong, 2007), which has been proved to increase the bullwhip effect

(Dejonckheere et al., 2003). Some authors have discussed other planning policies that can reduce the bullwhip effect (e.g., Boute et al., 2008; Dejonckheere et al., 2003; Gaalman, 2006; Gaalman & Disney, 2009; Graves, 1999). With an OUT policy, high target service levels have been proved to increase the bullwhip (Khosroshahi et al., 2016) and the same happens with long lead times (Agrawal et al., 2009; Chen et al., 2000a; Lee et al., 2000; Luong, 2007).

Some authors focused on analyses of the bullwhip effect generation when the final demand is identically and independently distributed (Dejonckheere et al., 2003; Kim et al., 2006). However, as demand autocorrelation has been proved to influence the bullwhip effect (Babai et al., 2016; Duc et al., 2008), many researchers have modelled the demand with ARIMA processes. In particular, the relationship between the bullwhip and an AR(1) final demand process has been widely studied (e.g., Chen et al., 2000a; Lee et al., 2000; Xu et al., 2001). More complex ARIMA models have also been taken into account: AR(p) models (Chandra & Grabis, 2005; Luong & Phien, 2007), ARMA(p,q) models (Alwan et al., 2003), ARIMA(0,1,1) models (Graves, 1999), and general ARIMA(p,d,q) models (Li et al., 2005). In all these studies, the stochastic process characterising the final demand and its parameters are assumed to be known. However, in real supply chains, only past demand observations are available and are used to infer the stochastic process and to estimate the parameters.

Some studies addressed the forecasting process to predict future demand. Efforts have been made to assess the impact of different forecasting methods on the bullwhip effect. The majority of the analytical studies assume an ARIMA framework for the demand process and an OUT policy, as discussed above. With these assumptions, the demand parameters are assumed to be known, some specific forecasting process is defined and a formulation for the bullwhip effect is given (Chandra & Grabis, 2005; Duc et al., 2008; Luong, 2007; Zhang, 2004). These works focus on the analysis of the relationship between the resulting bullwhip and its influencing factors (for instance, Luong (2007) considers an AR(1) demand and MMSE forecasts and focuses on the impact of the autoregressive parameter and of the lead time on the bullwhip effect). Other researchers assume some form of ARIMA demand process and do not tackle the problem of knowing the demand parameters as they use non-optimal forecasting methods such as Simple Moving Averages or Single Exponential Smoothing to predict the mean demand and forecast errors. In this case, the relationship between the bullwhip and the forecasting methods and their parameters are studied (Chen et al., 2000a; Dejonckheere et al., 2003; Xu et al., 2001). To the authors' knowledge, previous studies have not addressed the issue of unknown demand parameters, except for Hosoda & Disney (2009), which will be discussed later in this section.

The anti-bullwhip effect phenomenon (i.e., demand variability dampening) has been addressed by some authors (Alwan et al., 2003; Boute et al., 2008, 2014; Hosoda, 2005; Lee et al., 2000; Zhang, 2004) to analyse the impact of demand variability dampening in the supply chain. The variability of inventory has also been studied (Disney & Towill, 2003; Disney et al., 2004, 2006; Hosoda, 2005), as the inventory dynamics can affect the supply chain too. When considering inventory variability, sometimes demand bullwhip can be induced to enable inventory reductions (Disney et al., 2006). However, although the inventory variance issue has been proved to be relevant, it is beyond the scope of this paper.

All the cited papers, as discussed above, make strong assumptions about the final demand process. The demand process and the parameters of the process are assumed to be known. However, in real world supply chains, demand processes may be mis-specified (Hosoda & Disney, 2009) and, even if correctly specified, the parameters of the process are unknown and have to be estimated (Ali & Boylan, 2011). To the authors' knowledge, Hosoda & Disney (2009) are the only researchers to address the impact of demand process mis-specification. Their analysis shows that improving forecast accuracy does not always reduce supply chain costs. Thus, there is a need of further analytical investigations on the bullwhip effect in inventory systems characterised by uncertainty in the demand parameters. Consequently, this work focuses on the relationship between the bullwhip effect and uncertainty in the demand parameters.

3. Contributions of the paper

As discussed in Section 2, most theoretical bullwhip models are characterised by assumptions that make the models analytically tractable but hardly applicable in real contexts. The supply chains analysed through analytical models are characterised by a specific stochastic final demand with no uncertainty about the stochastic process or its parameters. Instead, in the real world, only a set of observations of the final demand model is available and, with this set, the stochastic process can be inferred and parameter values can be estimated.

To move towards more realistic analytical models, this work starts from the analytical framework analysed in the seminal work by Lee et al. (2000) and breaks one of its key assumptions. In Lee et al. (2000), a two echelon supply chain is investigated, whose inventories are planned through an order-up-to policy. The final demand follows a stationary autoregressive process $AR(1)$ (no uncertainty about the process) whose parameters are known (no uncertainty about the parameters). In this work, instead, the AR(1) process is assumed to be known (no uncertainty about the process), but the parameters of the process are assumed to be unknown and are estimated by an approximately unbiased estimator (uncertainty about the parameters). The demand is predicted through the Minimum Mean Squared Error forecast. As the results of this paper are compared with the outcomes of Lee et al. (2000), the impact of the uncertainty in the demand parameters can be appreciated.

The first contribution of this work is to show that not knowing the demand parameters has an impact on the bullwhip effect. The analysis also aims at understanding what demand and inventory characteristics amplify this effect. Particular attention is given to demonstrate that a negative (but unknown) autoregressive parameter can sometimes lead to a positive bullwhip effect for an AR(1) demand process, contrarily to the case of known demand parameters (Lee et al., 2000). The second contribution of this research is an investigation of the impact of the frequency of updating parameter estimates on the bullwhip effect. Results show that the more frequently the parameters are updated, the larger the bullwhip becomes. The result is consistent with the managerial insights of Hosoda & Disney (2009), which proved that, in two-stage supply chains, more accurate demand forecasts do not always lead to improved supply chain performance. However, the analyses in this paper differ from Hosoda & Disney (2009). They investigated the situation where an ARMA(1,1) process was mis-specified as an AR(1) process, and the frequency of forecast and parameter estimate updating was fixed. In this research, an AR(1) process is correctly specified but the parameters are estimated, and the estimates may be updated. Hosoda & Disney (2009) estimated parameters by minimizing the forecast errors, whereas this research uses an approximately unbiased estimator. Their analysis focuses on steady state behavior, whereas this paper investigates transient effects observed when there is limited demand data available. Hosoda & Disney (2009) focus on total supply chain costs but, in this paper, the bullwhip effect is evaluated, for direct comparison with the results from Lee et al. (2000). Minimisation of the bullwhip does not always guarantee minimisation of total supply chain costs (Disney et al., 2006); this interesting topic will be investigated in the next stage of this research.

The third contribution is to propose an analytical approximation to quantify the bullwhip effect in the case of unknown parameters of a known AR(1) process, when the demand forecasts are updated but the parameter estimates are not. A simulation model is used to show that the analytical approximation, in the case of unknown demand parameters, is able to predict the bullwhip more accurately than the one proposed in Lee et al. (2000), while in the case of known parameters they lead to the same results.

4. System Description

The reference system analysed in Lee et al. (2000) (and in this work) is a single-product two-echelon supply chain composed of one retailer and one manufacturer. The demand faced by the retailer is a stochastic $AR(1)$ process:

$$d_t = \tau + \rho d_{t-1} + \epsilon_t, \tag{1}$$

where d_t is the demand at time t, τ a constant parameter, ρ the autoregressive parameter, ϵ the model error term, and t the time index.

In Lee et al. (2000), the following assumptions are made:

(a) the demand process is known ($AR(1)$ stationary process as in equation (1)),

(b) the parameters τ and ρ of the process are known,

(c) ϵ_t is *iid* normally distributed with mean 0 and (known) variance σ^2.

However, in real contexts, the only available information about the final customer demand is the record of past orders that customers issued to the retailer. Hence, mis-specification in the demand process or in its parameters should be taken into account, especially when few observations are available. In this paper, assumption (b) is relaxed. It is assumed that the parameters τ and ρ of the process are unknown and estimated with an approximately unbiased estimator. The change in assumption (b) is clearly needed because the parameters of a demand process can never be known with perfect accuracy. In future work, the relaxation of assumption (a) will be addressed (as in Hosoda & Disney (2009)) but, in the meantime, it shall be retained, to isolate the effect of parameter estimation, even though process identification may sometimes be in error. The assumption of *iid* normally distributed demand in assumption (c) is more reasonable, especially for fast-moving products, and hence it is retained. The assumption of the approximately unbiased estimator will be discussed later in the paper.

The retailer faces the final customer demand d_t and replenishes the inventory each period by issuing orders to the manufacturer, which is characterised by a lead time l. The planning policy in use at the retailer echelon is the order-up-to policy with a review every period; thus, the out-of-control period of the retailer OUT policy is equal to $l + 1$ periods. The forecasting process must predict the future out-of-control period demand and the variability of forecast errors. Then, the planning policy uses the forecasting outputs to issue the replenishment order to the manufacturer.

4.1. Forecasting process

To predict demand in future periods after time t, information about demand up to time t is available but some estimates of the parameters of the process are needed. Let $\widehat{\tau}_t$ and $\widehat{\rho}_t$ be unbiased estimates of the parameters τ and ρ, respectively, calculated using the information about the demand up to and including time t. The forecasts of demand in periods $t+1, \ldots, t+l+1$ are needed as input in the planning process and they are denoted by the variables $\widehat{d}_{t+1}, \ldots, \widehat{d}_{t+l+1}$. These quantities are the estimates of the expected values of $d_{t+1}, \ldots, d_{t+l+1}$ respectively, using demand data up to and including time t. Thus, they can be calculated as in the following, exploiting the properties of AR(1) processes:

$$\widehat{d}_{t+k} = \widehat{\mathbf{E}}\left[d_{t+k}|d_t\right] = \widehat{\mathbf{E}}\left[\tau + \rho d_{t+k-1} + \epsilon_{t+k}|d_t\right] = \widehat{\tau}_t + \widehat{\rho}_t \widehat{d}_{t+k-1} \quad k = 1, \ldots, l+1. \tag{2}$$

Using these expressions, the forecast of the expected demand in the out-of-control period is:

$$\widehat{m}_{t|d_t} = \sum_{k=1}^{l+1} \widehat{d}_{t+k} = \frac{\widehat{\tau}_t}{1-\widehat{\rho}_t}\left\{l+1-\sum_{k=1}^{l+1}\widehat{\rho}_t^k\right\} + \frac{\widehat{\rho}_t\left(1-\widehat{\rho}_t^{l+1}\right)}{1-\widehat{\rho}_t} d_t. \tag{3}$$

The forecasts are assumed to be unbiased and the estimate of the variance of the forecast errors is denoted by $\widehat{v}_{t|d_t}$ and it is defined as the expected squared forecast error:

$$\widehat{v}_{t|d_t} = \mathbf{E}\left[\left(\sum_{k=1}^{l+1} d_{t+k} - \sum_{k=1}^{l+1}\widehat{d}_{t+k}\right)^2 \middle| d_t\right]. \tag{4}$$

4.2. Planning process

At the end of period t, the retailer places an order y_t to bring the inventory position up to the order-up-to level S_t, which is given by:

$$S_t = \widehat{m}_{t|d_t} + z_\alpha \sqrt{\widehat{v}_{t|d_t}}, \tag{5}$$

where $\widehat{m}_{t|d_t}$ is the estimated expected demand during the out-of-control period and $z_\alpha\sqrt{\widehat{v}_{t|d_t}}$ is the safety stock. Specifically, z_α is the α-quantile of the forecast error distribution, with α being the target Cycle Service Level, and $\widehat{v}_{t|d_t}$ the estimated variance of forecast errors. The terms $\widehat{m}_{t|d_t}$ and $\widehat{v}_{t|d_t}$ are conditioned on d_t as they are calculated using all the information about the demand up to and including time t (equations (3) and (4)). The replenishment order y_t issued to the manufacturer at time t is:

$$y_t = d_t + \left(S_t - S_{t-1}\right), \tag{6}$$

which is composed of the observed final demand d_t and the difference between the OUT levels at times t and $t-1$.

5. The impact of the estimation of unknown parameters on the bullwhip effect

As the difference between Lee et al. (2000) and the system studied in this paper is the uncertainty in demand parameters, a first analysis must quantify the effect of the parameter estimation on the bullwhip effect.

In the case of known demand parameters, the bullwhip effect (namely, $BE^{(K)}$) is derived in Lee et al. (2000) as:

$$BE^{(K)} = \frac{1-\rho^2}{(1-\rho)^2}\left(\frac{\rho^{2l+4}(1-\rho)}{1+\rho} + (1-\rho^{l+2})^2\right). \tag{7}$$

A simulation will be used to assess any differences between the bullwhip effect of an inventory system with unknown demand parameters and that predicted by equation (7).

5.1. The simulation model

The issue addressed by the simulation is to ascertain whether not knowing the true values of the demand parameters amplifies the bullwhip. Thus, the simulation is run to check the impact of the uncertainty in the demand parameters on the bullwhip effect by comparing the two systems previously described. The bullwhip characterising a system with known demand parameters can be predicted by the analytical formulation of equation (7), whereas there is currently no formulation available to characterise the bullwhip affecting a system with unknown parameters. Therefore, a simulation model is used to represent an inventory system with unknown parameters and to observe the magnitude of the bullwhip.

The simulation model represents the system with unknown demand parameters, as follows. At the beginning of the simulation, during the initialisation period (specified later in this section), the parameters of the demand process are estimated. Thereafter, the forecast and planning processes are simulated in each time period. Thus, at the end of each time t, the demand is observed, the forecast is updated and the replenishment order is issued. The parameters of the demand process are re-estimated with a frequency that varies within the experiments.

The estimator of the demand parameters used in the simulation is the linear-bias-correction (LBC) of the Ordinary Least Square (OLS) estimator (Kendall, 1954; MacKinnon & Smith, 1998). Given the Ordinary Least Squares estimator $\check{\rho}$, the LBC estimator $\hat{\rho}$ is calculated as:

$$\hat{\rho} = \frac{1}{n-3}(n\check{\rho}+1), \tag{8}$$

where n is the sample size, i.e., the number of the most recent observations used to compute the OLS estimator on each occasion the estimates are updated. As discussed later in the paper, this value varies between 12 and 48 observations in the simulation.

The OLS has been corrected according to equation (8) as the sample sizes used in the simulations are not large enough to assure the OLS to be unbiased. (It has been proved that with small sample sizes the OLS estimator is biased (MacKinnon & Smith, 1998)). Moreover, the bias has been proved to be approximately a linear function of ρ in the range $-0.85 \leq \rho \leq 0.85$ (MacKinnon & Smith, 1998). Hence, the linear-bias-correction is used as an approximately unbiased estimator.

Within the simulation, the demand is predicted by equation (3) and the variability of forecast errors is estimated with the Root Mean Square Error (RMSE). The RMSE is calculated considering all the historical errors available from the end of the initialisation period and, at each time, its value is updated by including the last error (e.g., when 10 historical forecast errors are available, all of them are used to calculate the RMSE; the next period, when 11 forecast are available, then all the 11 errors are used). Finally, the OUT level is calculated according to equation (5) and the replenishment order by using equation (6). At the end of the simulation, the bullwhip effect is calculated as the ratio of the variances of the observed demand series and the issued replenishment order series, to make the value directly comparable with Lee et al. (2000).

5.2. Experimental Design

The simulation experiments have been carried out by varying the values of some of the parameters related to the demand process, to the planning policy and to the estimation procedure. Table 1 summarises the factors investigated in the simulation experiment; all 4725 combinations are analysed.

The first two rows of the table are related to the true demand parameters. At the beginning of the simulation, the demand is generated according to the given AR(1) distribution. The parameter τ is set to 200 and never varied (as the bullwhip is independent of this parameter if the upstream and downstream demands have the same τ value). The autoregressive parameter ρ has been varied between negative and positive values, as it has been suggested that the variability of demand amplifies in the case of positive autocorrelation and smooths in the case of negative autocorrelation (Lee et al., 2000). The factor ρ has been limited to a minimum value $\rho = -0.6$ and to a maximum value $\rho = 0.6$, as previous empirical studies have shown that in real industries ρ usually does not reach magnitudes of 0.7 or greater (Ali et al., 2012; Erkip et al., 1990; Lee et al., 2000) (the estimator used in the analysis is

Table 1: The impact of the unknown parameter estimation: Experimental Design

Factor	**Notation**	**Range**
Autoregressive parameter	ρ	$\{-0.6, -0.3, 0, 0.3, 0.6\}$
Variability of the error term	σ	$\{1, 2, 5\}$
Sample size	n	$\{12, 24, 48\}$
Estimate updating interval	δ	$\{1, 2, 3, 8, 12, 52, \text{never}\}$
Lead time	l	$\{1, 2, 5\}$
Target Cycle Service Level	α	$\{60\%, 70\%, 80\%, 90\%, 99\%\}$

approximately unbiased for ρ values within the range $-0.85 \leq \rho \leq 0.85$). The standard deviation σ of the error term has been varied to check if the bullwhip is σ-independent also in the context of unknown demand parameters, as it happens in the case of known parameters (see equation (7)).

The third and fourth rows of Table 1 are related to the parameter estimation process. In the third row, the sample size n corresponds to the number of demand observations used to estimate the parameters at the beginning of the simulation and on every occasion that the parameters are re-estimated. According to the experimental design, the sample size n varies between 12 and 48 demand observations. This wide range has been chosen to encompass a realistic set of possible scenarios. The shorter sample sizes correspond to less mature products, enabling an examination of the sensitivity of results to short histories. The longer sample sizes correspond to more mature products.

In the fourth row, the estimate updating interval, δ, relates to how often the parameter estimates are updated. If δ takes a finite value, then the estimates are updated during the simulation (for instance, $\delta = 8$ means that the estimates are updated every 8 time periods), whereas if $\delta =$ never, then the parameters of the demand process are estimated at the beginning of the simulation and never updated. Updating the parameter values every δ time periods is the more realistic choice, as in reality having new demand observations usually leads to updating the estimates of the demand parameters; however, as the results of the simulation will show later in the paper, updating the estimates increases the bullwhip effect. Thus, to avoid the additional effect generated by updating the estimates, the case of never updating (i.e., $\delta =$ never) is included in the experimental design. Moreover, by analysing the simulations with a small n, it is possible to estimate the consequence on the bullwhip effect of less accurate estimates, without any possible interference of the effect of updating the estimate itself on the bullwhip.

Finally, the last two rows of Table 1 are related to the planning process. The lead time has been varied between 1 and 5 time periods (further experiments have been made with longer lead times, which led to similar results in terms of the relationship between the bullwhip and the lead time), and the target Cycle Service Level has been varied between 60% and 99%, to give an appreciation of the bullwhip effect over a wide range of realistic Cycle Service Level targets.

The experimental design is composed of 4725 combinations of factors, which have been all evaluated with simulations. For reasons of conciseness, the results of only some subsets are shown in the paper; however, exhaustive results are included in the Supplementary Material. For each combination of factors, $K = 1000$ replicates are performed. In each replicate, the simulation lasts $L = 1000 + n + l + 4$ time periods, where $n + l + 4$ is the initialisation period [1]. The bullwhip effect is measured over the last 1000 time periods. For each combination of factors, the single replicate bullwhip effect values are averaged over the 1000 replicates and the average value is compared to the bullwhip effect found by using equation (7).

5.3. Simulation results

As previously discussed, the objective of the experiment is to compare the bullwhip effect generated by an inventory system with unknown demand parameters, namely $BE^{(U)}$, and the one generated by an inventory system with known demand parameters, namely $BE^{(K)}$, to show the effect of the uncertainty of the demand parameters.

Before discussing the results, the assumptions made throughout the paper have been checked. First, $\widehat{\rho}$ is assumed to be unbiased. The results showed that, for all the combinations of factors, the mean error of $\widehat{\rho}$ is included in $[-0.07; 0.01]$. Also the forecasts are assumed to be unbiased and the results showed that, for all the combinations of factors, the MPE (mean percentage error) is included in $[-0.03\%; 0.03\%]$. Thus, the main assumptions are proved to hold within the simulations.

The results will show how the bullwhip effect changes when the values of the factors change. As will be discussed, the bullwhip is influenced by two phenomena: the uncertainty of the parameter values and

[1] In the initialisation period, the first time period is used to set the first observation of the final demand $\left(d_1 = \frac{\tau}{1-\rho}\right)$. Then, n demand observations are used to estimate $\widehat{\tau}$ and $\widehat{\rho}$. In the next time period, the first one-step ahead forecast is derived, after l time periods the first forecast error is calculated, and the last two time periods are used to have at least 2 observations to derive the RMSE of forecast errors.

the updating of their estimates. In the following, the two phenomena are assessed separately, starting with the effect of updating.

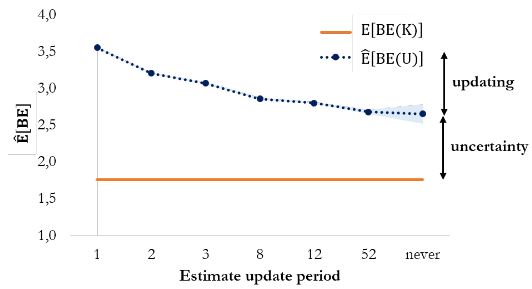

Figure 1: Influence of update interval (δ): reducing δ magnifies the bullwhip.

The effect of updating parameter estimates

Figure 1 shows how the bullwhip effect varies when the value of δ varies. It refers to the case in which: $\rho = 0.3, \sigma = 1, l = 1, n = 12, \alpha = 99\%$; however, similar results are obtained for each combination of factors, as shown in the Supplementary Material. In systems with known demand parameters, no estimate has to be made; thus the bullwhip is independent of the value of σ. The orange straight line, which shows the values of $BE^{(K)}$ calculated by equation (7), is constant. The blue dotted line shows the bullwhip in inventories with unknown demand parameters, i.e., the $BE^{(U)}$ average values. Each point of the dotted line is related to a specific value of δ, reported on the horizontal axis. As each point represents the average value over 1000 replicates, the light blue shadings show the 95%-confidence intervals for the mean values (the shadings are reported for each point; however, for some of them, the confidence interval is too small to be appreciated in the figure). The graph shows that the bullwhip generated in the simulation $BE^{(U)}$ decreases when δ increases. Reading the figure from the right to the left, if the parameters are never updated, the bullwhip is greater than that given by equation (7). As the only difference between the two systems is the assumption about the uncertainty in the parameters, this difference is caused by the effect of the *uncertainty* of the parameters of the demand distribution, which will be examined later in the section. If the parameter estimates are updated (δ from 1 to 52), then the bullwhip increases even more, and this marginal increase is the effect of the *updating* on the bullwhip. The more frequent the update is, the larger the bullwhip becomes. Thus, the results show that updating the parameters increases the bullwhip effect. Moreover, the results are consistent with the managerial insights of Hosoda & Disney (2009), showing that more accurate forecasts do not always imply improved supply chain performance.

The effect of uncertainty of demand estimates

To evaluate the effect of the uncertainty on the bullwhip, experiments with no updating (δ = never) have been performed (to isolate the impact of uncertainty from the impact of updating the parameters). The objective is to understand the causes of the difference between the bullwhip predicted by equation (7) and that measured in the simulations, for a range of sample sizes and for no updating.

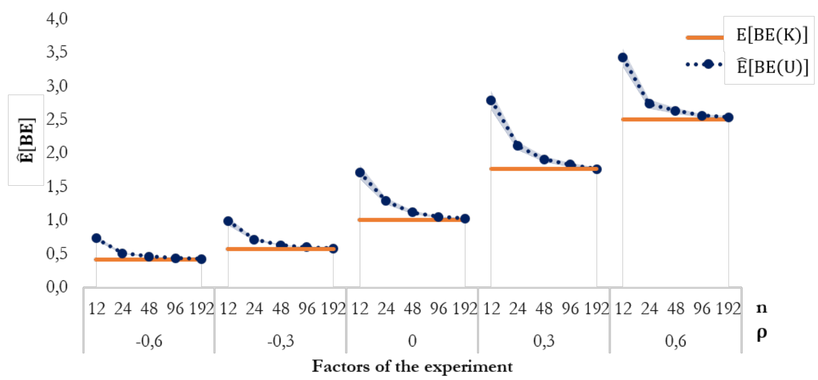

Figure 2: Influence of sample size (n) and autoregressive parameter (ρ): increasing n and ρ magnifies the bullwhip.

The effect of the uncertainty of the demand process parameters, and, hence, of the mis-specification of the model, can be appreciated by varying the sample size. The sample size n counts the number of observations used to estimate the parameters of the demand process. Figure 2 shows how the bullwhip varies with different values of n and ρ. The parameter ρ varies in the horizontal axis according to the experimental design. The range of values of n has been extended to include larger sample sizes ($n = 96, 192$). This enables the difference between the blue dotted lines ($BE^{(U)}$) and the orange ($BE^{(K)}$) to be appreciated when more demand observations become available. For all the values of the autoregressive parameter, the more demand observations are available to estimate the parameters, the lower the bullwhip. Moreover, with large values of n, there is virtually no difference between the two lines, meaning that equation (7) is a very good approximation of the bullwhip found in the simulation. With small values of n, instead, $BE^{(U)}$ and $BE^{(K)}$ are markedly different, and the effect of the uncertainty can be appreciated. When few demand observations are available, the estimates of the demand parameters are poor. In this case, the demand parameters cannot be assumed to be known and, hence, equation (7) becomes only a rough approximation for the bullwhip; indeed, equation (7) underestimates the bullwhip by a considerable margin in this case. On the other hand, in the case of a large sample size, the two lines converge. When a large number of demand observations is available, the estimates of the demand parameters are more accurate and closer to their true values.

The magnitude of the effect of the uncertainty on the bullwhip will be investigated in the remaining part of the section. The results are discussed in the following separately for each factor of the experiment except the updating interval, which is fixed as *never*.

The effect of n. As previously discussed, Figure 2 shows that the bullwhip generated in a system with unknown demand parameters can be larger than that generated with known demand parameters. Specifically, the smaller the sample size used to estimate the parameters, the larger the bullwhip effect. Also, the smaller the sample size, the larger the confidence interval of the average value of the bullwhip in the simulations (as shown later in Figure 3).

The effect of σ. In a system with known demand parameters, the bullwhip is invariant with respect to σ and n. This is demonstrated in the exact result of Lee et al. (2000), listed as equation (7) in this paper. For unknown demand parameters, it is clear that the bullwhip is not invariant with respect to n. To assess the effect of σ, an ANOVA test has been performed directly on the single replicate values of the bullwhip effect with unknown parameters, $BE^{(U)}$. The results showed that there is no significant difference among the means of $BE^{(U)}$ groups with different values of σ. Therefore, in the case of unknown parameters and no updating, there is insufficient evidence to reject the null hypothesis that the bullwhip effect is not influenced by σ.

The effect of l. Figure 3 confirms that the discussed results hold for lead times greater than one. The two graphs represent the bullwhip effect values for the autoregressive parameter values $\rho = 0$ and $\rho = 0.6$. In each graph, n and l vary, whereas the other two values are fixed at $\alpha = 99\%$ and $\sigma = 1$. In the graphs of Figure 3, the distance between the lines reduces when the sample size increases; thus fewer demand observations increase the bullwhip also for larger values of the lead time. In systems with known demand parameters, the $BE^{(K)}$ bullwhip increases with the lead time only for positively autocorrelated final demands ($\rho > 0$), whereas it is equal to 1 for all lead times when $\rho = 0$. However, in the case of unknown parameters, the bullwhip increases with the lead time also in the case of zero autoregressive parameter (blue dotted lines in Figure 3(a)). The same holds with positive ρ (Figure

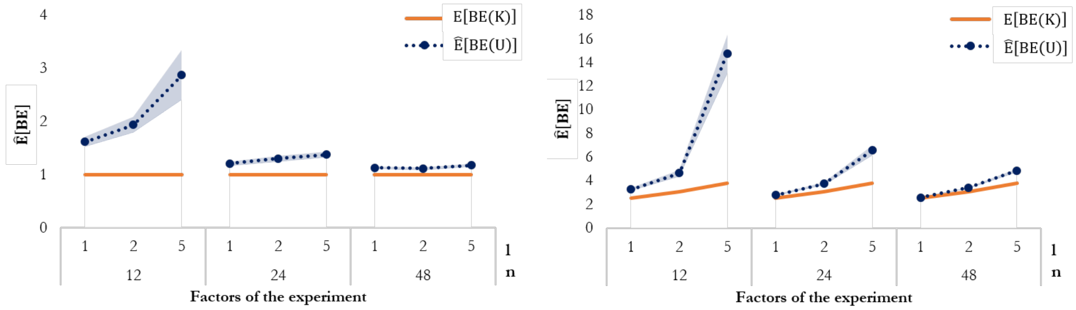

(a) $\rho = 0$ (b) $\rho = 0.6$

Figure 3: Influence of lead time (l): increasing the lead time magnifies the bullwhip.

3(b) shows the case $\rho = 0.6$ as an example). In this case, the orange line increases with lead time, but with a milder slope than the blue line.

The effect of α. The effect of different target Cycle Service Levels (α) has also been investigated. From equation (7), it is known that the bullwhip effect does not depend on α when the demand parameters are known. The same result seems to hold in the case of unknown parameters. The difference among the means of the $BE^{(U)}$ for different values of α has been found to be non-statistically significant.

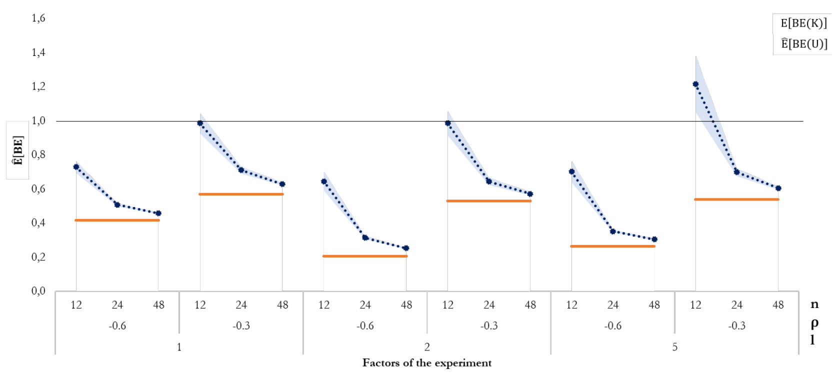

Figure 4: Influence of negative autoregressive parameter (ρ): a negative ρ does not always lead to an anti-bullwhip (see $n = 12$, $\rho = -0.3$, $l = 5$).

The effect of ρ. The last factor to be taken into account is the autoregressive parameter ρ. It is well established in the literature that, with known AR(1) demand process and known parameters, the demand variability is constant for $\rho = 0$, amplified for $\rho > 0$ and smoothed for $\rho < 0$ (Lee et al., 2000). Figure 3(a) shows a crucial situation: when the unknown demand process is iid ($\rho = 0$), but it is assumed to be autocorrelated (the true ρ value is unknown, and the autocorrelation is estimated by estimating the parameters of an AR(1) process), the bullwhip is larger than 1. Thus, the demand model mis-specification actually increases the bullwhip. Figure 3 show that the bullwhip $BE^{(U)}$ increases with ρ increasing also in systems with unknown demand parameters, with a larger magnitude than $BE^{(K)}$. Figure 4 shows the case of negative ρ (the other factors are set to $\sigma = 1$ and $\alpha = 99\%$). When the points are below 1, then there is an anti-bullwhip effect (i.e., the demand variability is smoothed rather than amplified). Interestingly, when $\rho = -0.3$ and $n = 12$, then $BE^{(U)}$ is greater than 1. This is an important result, as it shows that the variance of the demand is amplified by the effect of uncertainty of the parameters, which counterbalances the smoothing effect of having a negative autoregressive demand parameter. As a result, under unknown demand parameters, it is possible that there is a bullwhip effect, even when the autoregressive demand parameter is negative (while this never happens in systems with known demand parameters (Lee et al., 2000). Additional simulation experiments have been run to understand the ρ break-even value that separates bullwhip from anti-bullwhip (i.e., bullwhip greater and smaller than 1). For $\sigma = 1$, $\alpha = 99\%$, $l = 5$ and $n = 12$, the break-even ρ value has been found to lie between $\rho = -0.4$ and $\rho = -0.35$. The break-even value varies for each combination of factors, and so it must be estimated in each specific situation, when needed. From the managerial standpoint, this means that the set of the critical SKUs (i.e., the ones affected by bullwhip) is actually larger than the one that has been identified in the literature. Until now, the investigation of AR(1) process with a known negative autoregressive parameter has revealed an anti-bullwhip effect. The results in this paper show that this is not always true when ρ is unknown and needs to be estimated. Thus, SKUs that present negative autoregressive parameters can present some bullwhip too. Moreover, as for positive autoregressive parameters, the fewer the historical demand observations, the larger the bullwhip effect.

The skewness of $BE^{(U)}$ is positive for all the combinations of factors, meaning that the bullwhip values are concentrated within the left tail of the distribution.

In conclusion, the bullwhip affecting an inventory system characterised by unknown demand parameters tends to be larger than that affecting the same system but with known parameters. The increase of the bullwhip is even larger if there is a small sample size available for the estimation process. A new analytical approximation for $BE^{(U)}$ is proposed in the next section to to address the case of unknown demand parameters and to complement equation (7), which was designed for known parameters.

6. An approximate formula for the bullwhip effect with unknown parameters

As the comparison presented in the previous section showed some differences between the bullwhip effect generated in inventory systems with unknown and with known demand parameters, a new analytical approximation is here proposed for inventory systems characterised by unknown demand parameters and unbiased parameter estimators. The analytical approximation assumes that the estimates of the demand parameters do not change in time (i.e., $\delta =$ never).

In the same way as for equation (7), the bullwhip effect is calculated as the ratio of the variances of the replenishment orders (y_t) and the final demands (d_t). As d_t is an AR(1) process, its variance $var(d_t)$ is given by

$$var(d_t) = \frac{\sigma^2}{1-\rho^2}. \tag{9}$$

The variance $var(y_t)$ and an approximate formula for $BE^{(U)}$ are derived in the following sub-sections.

6.1. Variance of replenishment orders

The replenishment orders are calculated in the analysed system as in equation (6). Differently from the final demand, its variance is not known *a priori*, but must be calculated. The estimates of the demand parameters are assumed not to be updated over time (from now on, they are identified by $\widehat{\tau}$, $\widehat{\rho}$ instead of $\widehat{\tau}_t$, $\widehat{\rho}_t$).

Recalling equations (5) and (6), y_t has the following components:

$$y_t = d_t + \left(\widehat{m}_{t|d_t} + z_\alpha\sqrt{\widehat{v}_{t|d_t}}\right) - \left(\widehat{m}_{t-1|d_{t-1}} + z_\alpha\sqrt{\widehat{v}_{t-1|d_{t-1}}}\right).$$

Using the components in a different way, y_t can be written as:

$$y_t = d_t + \left(\widehat{m}_{t|d_t} - \widehat{m}_{t-1|d_{t-1}}\right) + z_\alpha\left(\sqrt{\widehat{v}_{t|d_t}} - \sqrt{\widehat{v}_{t-1|d_{t-1}}}\right). \tag{10}$$

Let Δm be the difference $\Delta m = \widehat{m}_{t|d_t} - \widehat{m}_{t-1|d_{t-1}}$ and let Δv be the difference $\Delta v = \sqrt{\widehat{v}_{t|d_t}} - \sqrt{\widehat{v}_{t-1|d_{t-1}}}$. Equation (10) can thus be rewritten as:

$$y_t = d_t + \Delta m + z_\alpha \Delta v. \tag{11}$$

Using the above defined components of y_t, its variance can be written as:

$$var(y_t) = var(d_t) + var(\Delta m) + z_\alpha^2 var(\Delta v) + 2cov(d_t, \Delta m) + 2z_\alpha cov(\Delta m, \Delta v) + 2z_\alpha cov(d_t, \Delta v). \tag{12}$$

Equation (12) shows that $var(y_t)$ depends on the variance of the final demand d_t, the variances of the differences of estimates Δm and Δv and the covariances among those terms, weighted by z_α where appropriate, which is the parameter related to the target Cycle Service Level α of the system. This equation is valid for all the inventory systems characterised by an order-up-to policy in which equation (6) holds, independently of the stochastic process of d_t. It is particularly relevant as it gives insights on what influences the variance of the replenishment orders: the final demand itself (d_t), the change in the forecast values that actually depends on the forecasting method (Δm), and the differences generated by updating the forecast errors (Δv).

All the terms of equation (12), except the ones related to Δv, can be expressed in compact formulae. The variance of the final demand $var(d_t)$, is already expressed in equation (9).
The variance $var(\Delta m)$ is given by:

$$var(\Delta m) = 2(1-\rho)\frac{\widehat{\rho}^2\left(1-\widehat{\rho}^{l+1}\right)^2}{\left(1-\widehat{\rho}\right)^2} var(d_t). \tag{13}$$

The covariance $cov(d_t, \Delta m)$ is given by:

$$cov(d_t, \Delta m) = (1-\rho)\frac{\widehat{\rho}\left(1-\widehat{\rho}^{l+1}\right)}{1-\widehat{\rho}} var(d_t). \tag{14}$$

The calculations of $var(\Delta m)$ and $cov(d_t, \Delta m)$ are provided in the Supplementary Material.
The terms related to Δv (i.e., $var(\Delta v)$, $cov(\Delta m, \Delta v)$ and $cov(d_t, \Delta v)$) are difficult to deal with in a compact formula. In the analytical approximation which follows, they have all been set to zero. This is equivalent to assuming that Δv does not affect the variability of the replenishment orders. This assumption makes the calculation of $var(y_t)$ only an approximation.

Summing up all the terms, the approximation of $var(y_t)$ becomes:

$$var(y_t) \sim var(d_t)\left[1 + 2(1-\rho)\frac{\widehat{\rho}\left(1-\widehat{\rho}^{l+1}\right)}{1-\widehat{\rho}}\left(1 + \frac{\widehat{\rho}\left(1-\widehat{\rho}^{l+1}\right)}{1-\widehat{\rho}}\right)\right]. \tag{15}$$

This approximation strictly depends on the assumptions previously made: (i) the parameter estimates are unbiased; (ii) the parameter estimates are not updated; (iii) updating the mean and variance estimates does not affect the variability of the replenishment orders (i.e., $var(\Delta v) = 0$, $cov(\Delta m, \Delta v) = 0$ and $cov(d_t, \Delta v) = 0$). The accuracy of this approximation will be assessed by means of a simulation experiment in Section 7.

6.2. The bullwhip effect

From equation (15), the bullwhip effect can be approximated by the following formula:

$$\tilde{BE}^{(U)} = 1 + 2(1-\rho)\frac{\widehat{\rho}\left(1-\widehat{\rho}^{l+1}\right)}{1-\widehat{\rho}}\left(1+\frac{\widehat{\rho}\left(1-\widehat{\rho}^{l+1}\right)}{1-\widehat{\rho}}\right). \tag{16}$$

Just as equation (15) approximates the true value of the variance of the replenishment demand, equation (16) approximates the bullwhip effect in the case of unknown demand parameters. Moreover, in the case of perfect estimation of the autoregressive parameter (i.e., $\widehat{\rho} = \rho$), equation (16) is equal to equation (7) (i.e., $\tilde{BE}^{(U)} = BE^{(K)}$). The proof of the equivalence is given in the Supplementary Material.

The form of the proposed approximation has some similarities with equation (7). As in equation (7), the bullwhip effect of systems with unknown demand parameters does not depend on σ while it does depend on the true autoregressive parameter ρ and on the lead time l. However, unlike equation (7), it also depends on the estimate of the autoregressive parameter, $\widehat{\rho}$. Thus, the bullwhip effect depends not only on a demand process parameter, but also on its estimate. The more accurate the estimate is, the closer $\tilde{BE}^{(U)}$ is to $BE^{(K)}$. This is strictly related to the influence of the sample size on the bullwhip effect previously discussed, which is not directly captured by the formula. The relationship between the bullwhip and the sample size will be analysed through simulations in Section 7.

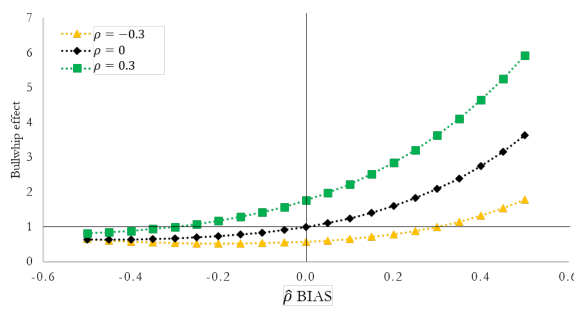

Figure 5: Relationship between the bullwhip and the bias of the estimate $\widehat{\rho}$: reading the figure from left to right, under-estimating ρ decreases the bullwhip effect, while over-estimating ρ increases it.

To appreciate the relationship between equations (7) and (16), consider the case of $\rho = 0$. When $\rho = 0$, according to equation (7), there is no bullwhip, i.e., $BE^{(K)} = 1$. Instead, as already shown in Section 5 (Figure 3(a)), $BE^{(U)}$ can exceed 1. This difference is explained by the dependence of $BE^{(U)}$ on the autoregressive parameter estimate $\widehat{\rho}$. Figure 5 displays the relationship between the bullwhip effect and the estimate $\widehat{\rho}$ with $l = 1$, and for $\rho = -0.3, 0, 0.3$. The position on the horizontal axis represents the bias of the estimate $\widehat{\rho}$, i.e., the value $\widehat{\rho}$ BIAS $= E[\widehat{\rho}] - \rho$, the vertical axis displays the bullwhip effect value, and each line collects the points for a specific ρ value. Thus, ρ is fixed for each curve, $\widehat{\rho}$ varies in the horizontal axis according to the $\widehat{\rho}$ BIAS variation, and each point is calculated from equation (16) according to the ρ of the curve and to the $\widehat{\rho}$ of the position on the horizontal axis. When $\rho = 0$ (black dotted curve) and the estimator measures it correctly (i.e., the point in the horizontal axis with $\widehat{\rho}$ BIAS $= 0$), then there is no bullwhip effect, exactly as equation (7) suggests. If the autoregressive parameter estimate is biased and, hence, different from 0, then the demand variability can be increased or smoothed, depending on the sign of the bias. According to Figure 5, an over-estimation of ρ (i.e., the points with $\widehat{\rho}$ BIAS > 0) increases the bullwhip effect, whereas an under-estimation of ρ (i.e., the points with $\widehat{\rho}$ BIAS < 0) decreases it. Interestingly, the effect of the wrong estimation of ρ on the bullwhip is not symmetrical, and this is true for all ρ values (see as example the green and yellow dotted curves in the figure): over-estimating the true value of ρ increases the bullwhip more than under-estimating ρ decreases it. This is of particular interest in the real world, when the presence of autocorrelation is not known *a priori* but must be estimated through data. In this case, for instance, even if the underlying process has no autocorrelation (hence there should not be any bullwhip), a mistake can be made in estimating the autoregressive parameter, and this can generate bullwhip (or smooth it). However, although an intentional under-estimation of ρ could lead to a lower bullwhip effect, it might have other effects on the inventory (such as, for instance, on the service level, on the investment in inventory and on the inventory variance). Moreover, particular attention should be paid to the line referring to a negative ρ value: in systems with known demand parameters, the demand variability is smoothed when the autoregressive parameter is negative. However, as the figure shows, a negative estimate of the autoregressive parameter does not automatically result in an anti-bullwhip when the parameter is unknown: for instance, when $\rho = -0.3$ and $\widehat{\rho} \geq 0.35$ (i.e., with $\widehat{\rho}$ BIAS ≥ 0.65), the final demand is amplified rather than smoothed.

7. A simulation to check the accuracy of the analytical approximation

The approximation proposed in Section 6 has been tested through a set of simulations. Specifically, the same simulation model described in Section 5.1 has been run. However, the aim here is to check the accuracy of the proposed analytical approximation, i.e., whether and to what extent it is close to the bullwhip effect measured in the simulation.

Within the experiment, $BE^{(K)}$ is evaluated by equation (7) for each combination of factors of the experimental design, the simulated $BE^{(U)}$ is evaluated from the final demand and the replenishment orders collected in the simulations, and the proposed analytical approximation $\tilde{BE}^{(U)}$ is evaluated in each single replicate by using equation (16) and the single-replicate value of $\hat{\rho}$. As for the $BE^{(U)}$, the single replicate $\tilde{BE}^{(U)}$ values are averaged over all the replicates for each combination of factors.

7.1. Experimental Design

The experimental design is based on that used in the previous experiment (Section 5.2, Table 1), but with some differences. In this experiment, the variability of the error term σ has been set equal to 1, as the analytical approximation of equation (16), as the formulation of $BE^{(K)}$, does not depend on σ. The estimate updating interval δ has been set to δ = never, as assumed in the analytical approximation. The final difference is that an additional factor has been included in the experimental design of the current simulation: the *performance collection length* CL. The factor CL is the time period in which the variability of the two series (demand and replenishment orders) are calculated. When CL is set to 24, for instance, then the simulation lasts $L = 24 + n + l + 4$ time periods (with $n + l + 4$ considered as the initialisation period, as in the previous experiment) and the variability of the final demand and of the replenishment orders is calculated over the last 24 time periods. In the previous experiment, CL was set to $CL = 1000$, whereas in this case it takes a value in the set $\{24, 48, 96, 1000\}$. The factor CL has been added to the list of factors to gain a better appreciation of the Δv component. Within the simulation, the variability of forecast errors is calculated with the RMSE of the forecast errors for windows of varying length. Thus, the larger the performance collection length, the more stable the variability of forecast errors $\hat{v}_{t|d_t}$ becomes (as a larger number of error observations leads to a more stable RMSE). Furthermore, the more stable the $\hat{v}_{t|d_t}$, the smaller becomes the magnitude of the difference $\Delta v = \sqrt{\hat{v}_{t|d_t}} - \sqrt{\hat{v}_{t-1|d_{t-1}}}$. Hence, with larger CL, the analytical approximation should be more accurate (as Δv should decrease).

The cases of small sample size n and large collection length CL are quite unrealistic, as if a long period is available to collect the demand, then it can be also available to estimate the demand parameters (leading, thus, to a large sample size). However, they are considered in the experimental design to gain a full understanding of the interaction of the factors.

All the other factors and parameters of the simulation are the same as in the previous experiment.

7.2. Simulation results

The main objectives of the simulation are: (i) to check whether equation (16) is a good approximation of the simulated bullwhip effect; (ii) to compare the proposed formulation with the one proposed in Lee et al. (2000) in a system with unknown demand parameters. More specifically, the aim of the section is (i) to compare $BE^{(U)}$ with $\tilde{BE}^{(U)}$ to check whether they are close to each other, and (ii) to analyse the pairwise comparisons $BE^{(U)}$–$BE^{(K)}$ and $BE^{(U)}$–$\tilde{BE}^{(U)}$ to check whether the simulated bullwhip $BE^{(U)}$ is better approximated by $BE^{(K)}$ (equation (7)) or $\tilde{BE}^{(U)}$ (equation (16)).

The accuracy of $BE^{(K)}$ and $\tilde{BE}^{(U)}$, with respect to $BE^{(U)}$, has been evaluated through error measures given by:

$$\text{Error } BE^{(K)} = \frac{BE^{(U)} - BE^{(K)}}{BE^{(U)}} \qquad \text{Error } \tilde{BE}^{(U)} = \frac{BE^{(U)} - \tilde{BE}^{(U)}}{BE^{(U)}}. \tag{17}$$

As in Section 5, the results are presented separately for each factor of the experiment.

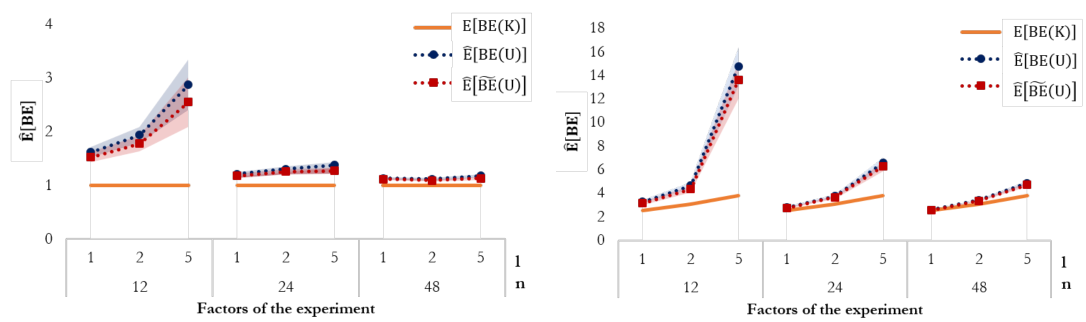

(a) $\rho = 0$

(b) $\rho = 0.6$

Figure 6: Simulation results (α = 99%, CL = 1000) on accuracy of new approximation: increasing the sample size (n) improves the accuracy; increasing the lead time (l) reduces the accuracy.

The effect of l. Figure 6 shows the results for α = 99%, CL = 1000, and lead time varying between 1 and 5 time periods. The figure is exactly the same as Figure 3 on which also a red dotted line, representing the values of equation (16), is reported. In the two graphs (only the cases of $\rho = 0$ and $\rho = 0.6$ are reported), the asymptotic bullwhip effect is shown (CL = 1000). In this case, the analytical approximation $\tilde{BE}^{(U)}$ is very close to the simulated $BE^{(U)}$. Both $BE^{(U)}$ and $\tilde{BE}^{(U)}$ increase when the lead time increases; however, the analytical approximation has a milder slope. Thus, the approximation is less accurate with larger lead times. Nevertheless, when the demand parameters are unknown, the proposed analytical approximation represents the asymptotic bullwip effect better than $BE^{(K)}$ (equation (7)). For the subset of combination of factors in Figure 6, $BE^{(K)}$ has an average error of 31%, whereas $\tilde{BE}^{(U)}$ has an average value of only 4% (the errors are positive as both equations underestimate the simulate bullwhip, but with a different magnitude). This quantifies the improvement in bullwhip accuracy achieved by the proposed approximation, which has been designed to catch the effect of unknown demand parameters. Furthermore, although the investigated lead time values are not greater than 5, further simulations have been made to check the cases of $l > 5$, and bullwhip accuracy improvements are also achieved in these cases.

Table 2: The accuracy of the analytical approximation with different values of CL.

Bullwhip effect	CL **values**			
	$CL = 24$	$CL = 48$	$CL = 96$	$CL = 1000$
Error $BE^{(K)}$	52%	39%	33%	28%
Error $\tilde{BE}^{(U)}$	38%	24%	14%	5%

The effect of CL. The effect of the factor CL is highlighted in Table 2. The table shows the errors of $BE^{(K)}$ and $\tilde{BE}^{(U)}$ for $\rho = 0.6$, $n = 12$, $l = 1$ and $\alpha = 99\%$. As previously mentioned, $BE^{(K)}$ does not change with different values of CL, but $BE^{(U)}$ decreases for larger CL.

Moreover, the difference between $BE^{(U)}$ and $\tilde{BE}^{(U)}$ becomes smaller when CL increases. Thus, the analytical approximation seems to accurately model the asymptotic behaviour of the bullwhip effect in the case of unknown parameters. However, when the analytical approximation to the bullwhip is calculated over a shorter collection length, the approximation is much less accurate, since it overlooks the instability of the estimates of the forecast errors. However, it still approximates the bullwhip better than equation (7) (the values in the second row of the table are smaller than those in the first row of Table 2).

The combined effect of the sample size n and the collection length CL is highlighted in Figure 7. The results confirm that the Lee et al. (2000) formula is a lower-bound for the simulated bullwhip effect when the parameters are unknown. It is a good lower-bound when the sample size is large but less accurate when the sample size is small. The analytical approximation is also a good lower-bound for large sample sizes; for smaller sample sizes it is less accurate, but a better lower-bound than the formula

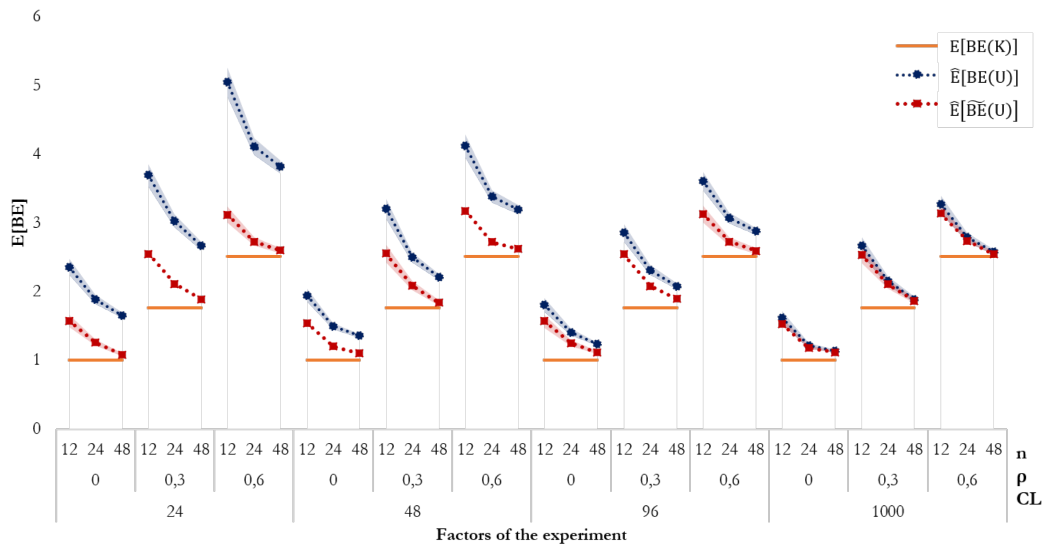

Figure 7: Combined effect of the sample size (n) and the collection length (CL): the new approximation is a lower bound of the simulated bullwhip and its accuracy improves as n and CL increase.

of Lee et al. (2000).

The analysis in Figure 7 shows the separate effects of the sample size and the collection length; in practical contexts these quantities would be usually the same. In Figure 7 the case of $n = CL = 48$ shows little benefit in using the analytical approximation, but the case of $n = CL = 24$ supports the earlier analysis showing the analytical approximation to be a better lower-bound.

The effect of α. Figure 8 shows the effect of the Cycle Service Level on the analytical approximation. In the figure, $l = 1$ and $\rho < 0$; however similar results are obtained for all the combinations of factors. Also, only the highest target Cycle Service Levels are shown ($\alpha = 90\%$ and $\alpha = 99\%$). From the figure, the difference of the accuracy of the proposed analytical approximation for different values of target Cycle Service Level α becomes clear. In Figure 8(a) (i.e., in the case of $\alpha = 99\%$), the red and the blue dotted lines are distant from one another for low values of n and CL. Instead, in Figure 8(b) ($\alpha = 90\%$), they are closer. This difference is explained by equation (12). As z_α multiplies all the terms of the equation that are assumed to be zero in the approximation $\tilde{BE}^{(U)}$ (because of the assumption $\Delta v = 0$), a small α reduces the effect of any violation of the assumption. As a consequence, with a smaller target Cycle Service Level α, the difference between $BE^{(U)}$ and $\tilde{BE}^{(U)}$ decreases. Specifically, in the simulations with $CL = 24$, the error of $\tilde{BE}^{(U)}$ is equal to 29.81% for $\alpha = 99\%$, 11.01% for $\alpha = 90\%$, and 0.35% for

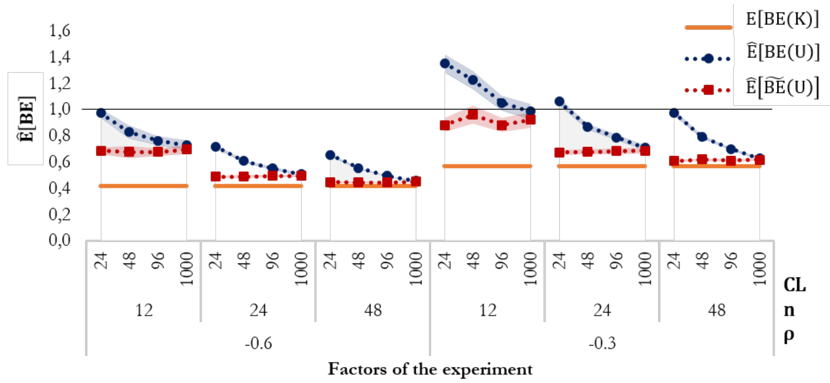

(a) $\alpha = 99\%$

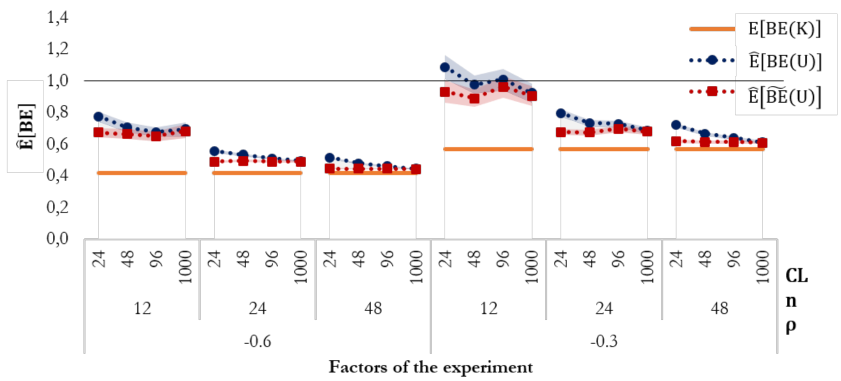

(b) $\alpha = 90\%$

Figure 8: Influence of the target Cycle Service Level (α): decreasing α (lower part of the figure) improves the accuracy of the approximation.

$\alpha = 60\%$ (the graph related to $\alpha = 60\%$ is not reported in the paper for reasons of conciseness). This confirms that the difference between the red and the blue lines is due to the assumption made for the analytical approximation.

The effect of ρ. Figures 6(a) and 7 show that, when $\rho = 0$, $\tilde{BE}^{(U)}$ is always greater than 1, meaning that the approximation is able to capture the effect that, in a system with unknown demand parameters, the bullwhip is amplified. Moreover, the accuracy of the analytical approximation does not change with negative autoregressive parameters. However, it is worth mentioning that although there are cases with $\rho < 0$ and $BE^{(U)} > 1$ (as previously discussed), the analytical approximation is not able to spot them.

For instance, Figure 8(a) shows that, when $l = 1$, $\rho = -0.3$, $n = 12$ and $CL = 24$, the approximation (red line) is less than 1 whereas the simulated bullwhip (blue line) is greater than 1. Regarding the bullwhip formulae, both $BE^{(K)}$ and $\tilde{BE}^{(U)}$ are below 1, without reaching the bullwhip region $BE^{(U)} > 1$. The same performance is shown in Figure 8(b), in the case of $\alpha = 90\%$. However, when comparing the accuracy of $\tilde{BE}^{(U)}$ and $BE^{(K)}$, $\tilde{BE}^{(U)}$ was found to be more accurate than $BE^{(K)}$.

In summary, when the bullwhip effect is calculated over long time periods, then the proposed analytical formulation $\tilde{BE}^{(U)}$ is a good approximation of the demand variability amplification. Moreover, comparing $\tilde{BE}^{(U)}$ and $BE^{(K)}$ as approximations for the bullwhip $BE(U)$ in the case of unknown demand parameters, $\tilde{BE}^{(U)}$ was found to be more accurate than $BE^{(K)}$. This is always true in the investigated cases, even when there are only few demand observations available to calculate the bullwhip effect (i.e., when CL takes small values). To conclude, as both $\tilde{BE}^{(U)}$ and $BE^{(K)}$ always underestimate $BE^{(U)}$ and $\tilde{BE}^{(U)}$ is always the closer approximation to the $BE^{(U)}$, then $\tilde{BE}^{(U)}$ can be identified as the better approximation for the bullwhip effect in systems with known AR(1) demand process and unknown demand parameters.

8. Conclusions, limitations and future research

The bullwhip effect is a very important issue for supply chains. For this reason, it has been widely studied in the literature. Specifically, this problem has been tackled both from a theoretical and an empirical point of view. From a theoretical standpoint, analytical models are usually proposed to identify the magnitude and the possible causes of the bullwhip effect. To reach this objective, various assumptions are made about the inventory system or the demand process. This generally makes analytical frameworks very far from being realistic and applicable in the real world.

The aim of this paper is to take a step forward in bridging the gap between the existing analytical models of the bullwhip effect and its effect on real supply chains. The proposed analytical framework starts from the seminal contribution of Lee et al. (2000), which considered a two-echelon supply chain with final demand following an autoregressive AR(1) process with known demand parameters. In the real world, organisations do not know the demand process or its parameters, and so the analytical expression for the bullwhip effect could be quite inaccurate. To partially address this issue, the assumption of knowing the parameters of the demand process is relaxed in the paper. Hence, differently from Lee et al. (2000), the final demand of the inventory system studied in the paper is characterised by unknown parameters, which are estimated to predict future demand.

There are three main contributions of this work, all of which relate to known AR(1) demand processes and Order-Up-To (OUT) inventory systems: 1) an evaluation of the effect of parameters being unknown on the bullwhip; 2) an evaluation of the effect of updating the parameter estimates; 3) the proposal of a new analytical approximation of the bullwhip effect in inventory systems with unknown demand parameters.

Regarding the first contribution, not knowing the parameters of the demand process has an effect on the bullwhip. Moreover, the difference is even more sizeable when few demand observations are available for the estimation process (and the uncertainty in the parameters is larger). Also, a negative autoregressive parameter does not always imply an anti-bullwhip effect, when it is unknown and, hence, estimated.

About the second contribution, the results show that updating the estimates more frequently (thus having more accurate forecasts) increases the bullwhip effect (and the related costs). This confirms the managerial insight of Hosoda & Disney (2009) that improved forecast accuracy does not necessarily imply improved supply chain performance.

In terms of the third contribution, the proposed formula is valid for the cases of no update of the demand parameters. It still depends on the true autocorrelation of the demand and on the lead time, as in the formula by Lee et al. (2000); however, it is also a function of the estimates of the autoregressive parameter. With this formulation, (i) the influence of not knowing the parameters has been analytically demonstrated and (ii) this influence has been quantified. The accuracy of the approximation has been evaluated through a simulation experiment. The results showed that both the proposed analytical approximation and the one proposed in Lee et al. (2000) underestimate the bullwhip effect in systems with unknown demand parameters. However, the proposed approximation is always more accurate than the one in Lee et al. (2000), as it also depends on the autoregressive parameter estimate. This is particularly true when few demand observations are available to estimate the demand parameters, and when the target Cycle Service Level is low.

The paper focuses on the impact of uncertainty of the demand parameter values on the bullwhip. Thus, this work has shown that in the real world, where the demand parameters are unknown, the generated bullwhip effect is higher than the one identified for inventory systems with known demand parameters. These results can give some insights from a managerial standpoint. Companies should be aware that the uncertainty has a tangible effect in the upstream demand process generation and, thus, it should be managed. Usually, companies do not check whether the fitted process or its parameters keep well fitting the observed demand over time. An implication of the work presented in this paper is that using the same estimates for the parameters while their values are changing can increase the bullwhip effect. However, frequent updating of the estimates can also increase the bullwhip if the true demand parameters are unchanging. This work has also addressed some practical issues. First, by showing that in some cases a negative autoregressive parameter does not always imply demand variability dampening, this work suggests to practitioners that the set of SKUs affected by bullwhip is broader than that suggested so far in the literature. The approximation can be considered a lower-bound of the real bullwhip affecting inventory systems.

The current research still has some limitations. In this paper, only the effect on the uncertainty of the parameters on the bullwhip effect was investigated and, to this aim, the demand process was assumed to be known. Although this assumption was convenient for analytical derivations, this limitation must be overcome in future research. Also, the effect of updating the estimates and of the frequency of the updates still has to be understood more deeply and quantified from an analytical perspective. The results of the paper have already shown that updating the parameter estimates can increase the bullwhip if the parameters are unchanging. However, further research is needed on how to choose the frequency of updating when the (true) parameters are changing over time. Furthermore, as the current analytical formula is an approximation, future research will be devoted to the search of an exact (or less approximate) formulation. To reach this goal, also the variability of the forecast errors must be analytically investigated and included. Although this work is able to give some insights to practitioners, further research will be devoted to empirically testing the proposed approximation.

References

Agrawal, S., Sengupta, R. N., & Shanker, K. (2009). Impact of information sharing and lead time on bullwhip effect and on-hand inventory. *European Journal of Operational Research*, *192*, 576–593.

Ali, M. M., & Boylan, J. E. (2011). Feasibility principles for downstream demand inference in supply chains. *Journal of the Operational Research Society*, *62*, 474–482.

Ali, M. M., Boylan, J. E., & Syntetos, A. A. (2012). Forecast errors and inventory performance under forecast information sharing. *International Journal of Forecasting*, *28*, 830–841.

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

Anderson, E. G., Fine, C. H., & Parker, G. G. (2000). Upstream volatility in the supply chain: The machine tool industry as a case study. *Production and Operations Management*, *9*, 239–261.

Babai, M. Z., Boylan, J. E., Syntetos, A. A., & Ali, M. M. (2016). Reduction of the value of information sharing as demand becomes strongly auto-correlated. *International Journal of Production Economics*, *181*, 130–135.

Blanchard, O. J. (1983). The production and inventory behavior of the American automobile industry. *Journal of Political Economy*, *91*, 365–400.

Boute, R. N., Disney, S. M., Lambrecht, M. R., & Van Houdt, B. (2008). A win-win solution for the bullwhip problem. *Production Planning and Control*, *19*, 702–711.

Boute, R. N., Disney, S. M., Lambrecht, M. R., & Van Houdt, B. (2014). Coordinating lead times and safety stocks under autocorrelated demand. *European Journal of Operational Research*, *232*, 52–63.

Bray, R. L., & Mendelson, H. (2015). Production smoothing and the bullwhip effect. *Manufacturing & Service Operations Management*, *17*, 208–220.

Chandra, C., & Grabis, J. (2005). Application of multi-steps forecasting for restraining the bullwhip effect and improving inventory performance under autoregressive demand. *European Journal of Operational Research*, *166*, 337–350.

Chen, F., Drezner, Z., Ryan, J. K., & Simchi-Levi, D. (2000a). Quantifying the bullwhip effect in a simple supply chain: The impact of forecasting, lead times, and information. *Management Science*, *46*, 436–443.

Chen, F., Ryan, J. K., & Simchi-Levi, D. (2000b). The impact of exponential smoothing forecasts on the bullwhip effect. *Naval Research Logistics (NRL)*, *47*, 269–286.

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

Disney, S. M. (2007). The dynamics of material flows in supply chains. *Available at SSRN 1013326*, .

Disney, S. M., Farasyn, I., Lambrecht, M., Towill, D. R., & Van de Velde, W. (2006). Taming the bullwhip effect whilst watching customer service in a single supply chain echelon. *European Journal of Operational Research*, *173*, 151–172.

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

Disney, S. M., Towill, D. R., & Van de Velde, W. (2004). Variance amplification and the golden ratio in production and inventory control. *International Journal of Production Economics*, *90*, 295–309.

Duc, T. T. H., Luong, H. T., & Kim, Y.-D. (2008). A measure of the bullwhip effect in supply chains with stochastic lead time. *The International Journal of Advanced Manufacturing Technology*, *38*, 1201–1212.

Edgehill, J., Olsmats, C., & Towill, D. (1987). Industrial case study of the dynamics of a close-coupled production-distribution system. In *Proceedings of the 9th International Conference on Production Research* (p. 2407).

Erkip, N., Hausman, W. H., & Nahmias, S. (1990). Optimal centralized ordering policies in multi-echelon inventory systems with correlated demands. *Management Science*, *36*, 381–392.

Forrester, J. W. (1958). Industrial dynamics: a major breakthrough for decision makers. *Harvard Business Review*, *36*, 37–66.

Gaalman, G. (2006). Bullwhip reduction for ARMA demand: The proportional order-up-to policy versus the full-state-feedback policy. *Automatica*, *42*, 1283–1290.

Gaalman, G., & Disney, S. M. (2009). On bullwhip in a family of order-up-to policies with ARMA (2, 2) demand and arbitrary lead-times. *International Journal of Production Economics*, *121*, 454–463.

Gill, P., & Abend, J. (1997). Wal-mart: The supply chain heavyweight champ. *Supply Chain Management Review*, *1*, 8–16.

Graves, S. C. (1999). A single-item inventory model for a nonstationary demand process. *Manufacturing & Service Operations Management*, *1*, 50–61.

Hammond, J. H. (1994). *Barilla spa*. Harvard Business School Boston, MA.

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

Hosoda, T., & Disney, S. M. (2009). Impact of market demand mis-specification on a two-level supply chain. *International Journal of Production Economics*, *121*, 739–751.

Hosoda, T., Naim, M. M., Disney, S. M., & Potter, A. (2008). Is there a benefit to sharing market sales information? Linking theory and practice. *Computers & Industrial Engineering*, *54*, 315–326.

Kendall, M. G. (1954). Note on bias in the estimation of autocorrelation. *Biometrika*, *41*, 403–404.

Khosroshahi, H., Husseini, S. M., & Marjani, M. (2016). The bullwhip effect in a 3-stage supply chain considering multiple retailers using a moving average method for demand forecasting. *Applied Mathematical Modelling*, *40*, 8934–8951.

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

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

Lee, H. L., So, K. C., & Tang, C. S. (2000). The value of information sharing in a two-level supply chain. *Management Science*, *46*, 626–643.

Lee, H. L., & Whang, S. (2000). Information sharing in a supply chain. *International Journal of Manufacturing Technology and Management*, *1*, 79–93.

Li, G., Wang, S., Yan, H., & Yu, G. (2005). Information transformation in a supply chain: a simulation study. *Computers & Operations Research*, *32*, 707–725.

Luong, H. T. (2007). Measure of bullwhip effect in supply chains with autoregressive demand process. *European Journal of Operational Research*, *180*, 1086–1097.

Luong, H. T., & Phien, N. H. (2007). Measure of bullwhip effect in supply chains: The case of high order autoregressive demand process. *European Journal of Operational Research*, *183*, 197–209.

MacKinnon, J. G., & Smith, A. A. (1998). Approximate bias correction in econometrics. *Journal of Econometrics*, *85*, 205–230.

Pastore, E., Alfieri, A., & Zotteri, G. (2019). An empirical investigation on the antecedents of the bullwhip effect: Evidence from the spare parts industry. *International Journal of Production Economics*, *209*, 121–133.

Rong, Y., Snyder, L. V., & Shen, Z.-J. M. (2008). Bullwhip and reverse bullwhip effects under the rationing game. *Available at SSRN 1240173*, .

Terwiesch, C., Ren, Z. J., Ho, T. H., & Cohen, M. A. (2005). An empirical analysis of forecast sharing in the semiconductor equipment supply chain. *Management Science*, *51*, 208–220.

Xu, K., Dong, Y., & Evers, P. T. (2001). Towards better coordination of the supply chain. *Transportation Research Part E: Logistics and Transportation Review*, *37*, 35–54.

Zhang, X. (2004). The impact of forecasting methods on the bullwhip effect. *International Journal of Production Economics*, *88*, 15–27.