MANAGEMENT SCIENCE
Vol. 55, No. 5, May 2009, pp. 781–797
ISSN 0025-1909 | EISSN 1526-5501 | 09 | 5505 | 0781
DOI 10.1287/mnsc.1080.0983

Information Sharing and Order Variability Control Under a Generalized Demand Model

Li Chen
Fuqua School of Business, Duke University, Durham, North Carolina 27708,
li.chen@duke.edu

Hau L. Lee
Graduate School of Business, Stanford University, Stanford, California 94305,
haulee@stanford.edu

The value of information sharing and how it could address the bullwhip effect have been the subject of studies in the literature. Most of these studies used different forms of demand models, assuming that no order smoothing was used by the retailer and that the supplier has full knowledge of the retailer's demand model and order policy. In this paper, we contribute to the literature by starting with a most general demand model, coupled with a smoothing policy for order variability control. In addition, we do not require that the supplier has full knowledge of the retailer's demand model and order policy, but instead let the retailer share its projected future orders (and freely revise them as the retailer sees fit). Under such a setting, we first obtain a unifying formula for the magnitude of the bullwhip effect. The formula indicates that it is the forecast correlation over the exposure period as a whole that determines the magnitude of the bullwhip effect. We then quantify the value of information sharing and generalize the existing results in the literature. Finally, we explore the optimal smoothing parameters that could benefit the total supply chain. The resulting optimal policy resembles the postponement strategy. We find that information sharing together with order postponement improves the supply chain performance, even though the order variability may amplify in some cases.

Key words: supply chain management; information sharing; inventory control; order smoothing; variability reduction; MMFE model
History: Received October 14, 2007; accepted December 4, 2008, by Ananth Iyer, operations and supply chain management. Published online in *Articles in Advance* February 19, 2009.

1. Introduction

Information sharing between supply chain partners has been viewed as one of the major means to improve the performance of the supply chain (e.g., see Kurt Salman Associates 1993). A retailer sharing point-of-sales (POS) data with a supplier has been considered one means to address the adverse problems of the bullwhip effect (Lee et al. 1997), and there has been an active stream of research on the value of information sharing in the presence of the bullwhip effect. This stream of research is often based on slightly more complex demand models (such as the integrated moving average IMA(0, 1, 1) in Graves 1999 and the autoregressive AR(1) in Lee et al. 2000), as otherwise the value of information sharing with independent and identically distributed (i.i.d.) demands is limited.

There are three interesting questions that arise in the literature of information sharing. First, what kind of value in information sharing can we obtain when demands are not i.i.d., IMA(0, 1, 1), or AR(1)? There were some mixed results on the value of information sharing in the literature. Are these results a function of the demand model used? Second, although most research indicated that the supplier would gain by having POS information, the supplier has to know the customer demand model as well as the underlying order policy to fully capture the value of information sharing. What happens if the supplier does not have the full knowledge of the underlying demand model and order policy? Third, most of this literature assumed that the retailer would pass on the orders to the supplier based on its demands, without any smoothing of the orders. But there seems to be evidence that some order smoothing exists in real life, which tends to reduce the bullwhip effect. When that is the case, what is the value of information sharing, and what should be the best smoothing policy for a retailer to use? This paper seeks to address these three questions.

There is a vast literature on the subject of information sharing and the bullwhip effect (see Chen 2003 for a comprehensive review). An empirical study by Clark and Hammond (1997) revealed that sales/forecast sharing via electronic data interchange, when coupled with the continuous replenishment process (a program similar in nature to vendor-managed inventory), can generate dramatic improvements in inventory turns for suppliers (e.g., Procter & Gamble and Campbell Soup in their case study). Bourland et al. (1996) studied the value of timely demand information when the retailer's order cycle and the supplier's production cycle are out of phase. Gavirneni et al. (1999) carried out a simulation study on the value of information in capacitated supply chains with i.i.d. demands. Cachon and Fisher (2000) studied the case with one supplier, N identical retailers, and i.i.d. demands. Graves (1999) quantified the bullwhip effect under the integrated moving average IMA$(0,1,1)$ demand model and showed that there are no additional benefits to sharing downstream POS information. Lee et al. (2000) showed that under the AR(1) demand model, sharing POS information could provide significant value to the supplier, with the value being greater when the autocorrelation is greater. Under the same model, however, Raghunathan (2001) showed that if the supplier knows the parameters of the retailer's demand model and order policy and is able to utilize all the order history information from the retailer, then there is little incremental value to the supplier when POS data were shared. Gaur et al. (2005) further extended the result of Raghunathan to a general autoregressive moving average (ARMA) model and showed the value of information sharing when the retailer's demand is not inferrable from the order history.

Through a moving average demand forecast method, Chen et al. (2000) demonstrated that the bullwhip effect can be reduced by centralizing demand information. Aviv (2003) used a general time-series framework to study the benefits of information sharing in supply chain management. Miyaoka and Hausman (2004) showed that under the IMA$(0,1,1)$ demand model, using "stale" forecasts can reduce the bullwhip effect and improve fulfillment from the supplier to the retailer.

We note that in modeling the value of information sharing, most studies assume that the supplier has full knowledge of the underlying demand model and the order policy used by the retailer. In reality, this is an aggressive assumption. When the supplier does not have such knowledge, the value of POS or the historical order data may be limited. Instead, it may be more useful for the supplier to know of the retailer's forecast and/or some form of advance demand information. One such form of advanced demand information is the retailer's projected future orders (but of course, the retailer is not bound by such projections and could revise the projections over time). In an empirical study, Terwiesch et al. (2005) reported that sharing projected orders is a common practice in a semiconductor equipment supply chain. Although this practice is used for the purpose of capacity reservation according to their study, we will show in this paper that sharing projected future orders is also an effective way to pass demand information along a supply chain when capacity is not a major concern.

Some companies may go even further to collaborate on forecasting. For instance, the *collaborative, planning, forecasting, and replenishment* initiative started by the Voluntary Interindustry Commerce Solutions Association (1999) has reported early successes from forecast reconciliation between retailers and suppliers. The benefits of such collaborative forecasting have been modeled by Aviv (2001, 2002, 2007) using the i.i.d. and AR(1) demand models with additional explanatory variables that represent the market signals observed by the supply chain partners.

In a recent study, Cachon et al. (2007) examined the industry-level bullwhip effect using empirical data from the U.S. Census Bureau. They concluded that, at the industry level, the production-smoothing effect is more pronounced than the variability amplification effect, which suggests the potential existence of a form of variability control in supply chains—the orders from the retailer have been smoothed in some way to reduce variability. Order variability could be a legitimate cost concern to the retailer's ordering decision. For example, the retailer may incur additional shipping and handling capacity cost when its order exceeds the normal variation range (Balakrishnan et al. 2004).

Order variability control has also been studied in the literature. Graves et al. (1998) studied the production-smoothing policies for a single production stage to optimize the trade-off between production capacity and inventory. Balakrishnan et al. (2004) studied the benefit of a convex-combinational order-smoothing policy under the i.i.d. demand model. Liu et al. (2005) proposed a near-optimal policy to reduce order variability for the AR(1) demand model by imposing a constraint on the maximum allowable variability level. Aviv (2007) proposed an adaptive linear quadratic Gaussian production policy to minimize a quadratic cost function that captures inventory considerations, production smoothing, and adherence-to-plans.

In this paper, we contribute to the literature by starting with a most general demand model, coupled with a smoothing policy for order variability control. We explore the optimal smoothing parameters that could benefit the total supply chain. In addition, we also do not require that the supplier has full knowledge of the retailer's demand model and order policy; instead we let the retailer share its projected future orders (and freely revise them as the retailer sees fit). However, we do assume that the supplier will use the historical projected order information to determine the stochastic characterization of the order revision process, which we believe is a quite reasonable assumption, as it is what the supplier would naturally do. Under such a setting, we explore the value of information sharing by the retailer and the resulting bullwhip effect.

The pioneering work done by Aviv (2001, 2002, 2003, 2007) in studying information sharing and also delineating the difference between uncertainty and variability is particularly relevant to this paper. It is important to point out here the major differences between Aviv's work and ours. First, the general demand model used in Aviv's work is the linear state-space demand model (see Aviv 2003). As we will illustrate in §2, the linear state-space model, like the ARMA time-series model, is a special case of our generalized demand model. Because our model does not rely on specific linear state-space structure, the derived results are more concise and general. Second, the retailer's ordering policy studied in Aviv's work is a special case of the generalized order-up-to policy proposed by our paper. Our generalized order policy enables us to search for the optimal order-smoothing weight to minimize the total supply chain costs. Third, Aviv (2003) considered the case when the supplier has no knowledge of the retailer's demand model and ordering policy. He showed that in such a case, the retailer's order process can be characterized as an AR(∞) process, which is virtually impossible to estimate, as it involves an infinite number of model parameters. In contrast, we show a stronger result that in such case, having the retailer share the projected future orders would achieve the full value of information; the supplier only needs to determine the variance-covariance structure of the order revision process based on historical order information, which is very easy to do. Finally, Aviv (2003) presented a method for estimating the necessary safety stock when the system is not decoupled. In our paper, to derive close-form results, we only consider the simplified case when the system is decoupled by assuming the existence of a secondary source for supply. But we argue that this is not a serious restriction to the insights derived from our analysis because many retailers in real world, such as Wal-Mart, generally require suppliers to achieve a 95%–98% service level. Given such a high service-level agreement (which is, of course, not necessarily optimal from the integrated system perspective), the performance of an integrated system can be fairly well approximated by a decoupled system.

The generalized demand model that we use in the paper is the Martingale model of forecast evolution (MMFE). The MMFE model was first introduced as a demand model for inventory research by Hausman (1969), Graves et al. (1986), and Heath and Jackson (1994). In a field study with Eastman Kodak, Graves et al. (1998) validated the MMFE model with empirical sales and forecast data. Aviv (2001, 2002, 2007) used the MMFE framework to model the evolution of the market signals, or forecast explanatory power, available to the supply chain partners under the collaborative forecasting initiative. The advance demand information model by Gallego and Ozer (2001) is essentially an MMFE model. Most time-series demand models can be interpreted as a special case of the MMFE model, such as the AR(1) model (Lee et al. 2000, Raghunathan 2001), the IMA(0, 1, 1) model (Graves 1999, Miyaoka and Hausman 2004), the general ARMA model (Gilbert 2005, Gaur et al. 2005), and the linear state-space demand model (Aviv 2003). Other papers that make use of the MMFE demand model include Gullu (1996, 1997), Toktay and Wein (2001), Iida and Zipkin (2006), and Lu et al. (2006).

The remainder of this paper is organized as follows. Section 2 introduces the general MMFE demand model. In §3, we characterize a generalized order-up-to policy. In §4, we quantify the benefits of information sharing under the generalized order-up-to policy. We derive the optimal order variability control in §5. Section 6 contains extensions of the current model framework and our concluding remarks. All proofs are presented in the appendix.

2. A Generalized Demand Model

Consider a supply chain that consists of a retailer and a supplier. The external demand faced by the retailer is unfolded according to the MMFE process (Hausman 1969, Graves et al. 1986, Heath and Jackson 1994). Specifically, we assume that the demand for a period t is given by

$$D_t = \mu + \sum_{i=0}^{\infty} \epsilon_{t-i,\,t}, \tag{1}$$

where $\epsilon_{t-i,\,t}$ is the incremental information obtained in period $t-i$ with respect to demand D_t. For all $i \geq 0$, $\epsilon_{t-i,\,t}$ is mutually independent, stationary, and normally distributed with $N(0, \sigma_i^2)$. Let $\sigma^2 = \sum_{i=0}^{\infty} \sigma_i^2$. If $\sigma^2 < \infty$, then the demand in each period is stationary, with a normal distribution $N(\mu, \sigma^2)$ (which may be correlated across different periods). If $\sigma^2 = \infty$, then only the conditional variance, say, from period 0—i.e., $\text{var}(D_t \mid \epsilon_{t-i,\,t}, i \geq t) = \sum_{i=0}^{t-1} \sigma_i^2 < \infty$—is a meaningful measure. For ease of exposition, we will assume $\sigma^2 < \infty$ throughout the paper; in the case of $\sigma^2 = \infty$, the results obtained in this paper should be modified with the conditional expectation.

We note that the variance σ^2 reflects the total variability of demand. It does not necessarily represent the demand uncertainty faced by the retailer because the uncertainty can be partially resolved by the incremental demand information $\epsilon_{t-i,\,t}$ as time goes by. The

Table 1 Illustration of the MMFE Process

	Forecast revision for future demands in each period							
Demand	$\cdots$	0	1	2	3	4	5	6 $\cdots$
$D_0=\mu$	$+\cdots$	$+\ \epsilon_0$						
$D_1=\mu$	$+\cdots$	$+\ \epsilon_{0,1}$	$+\ \epsilon_1$					
$D_2=\mu$	$+\cdots$	$+\ \epsilon_{0,2}$	$+\ \epsilon_{1,2}$	$+\ \epsilon_2$				
$D_3=\mu$	$+\cdots$	$+\ \epsilon_{0,3}$	$+\ \epsilon_{1,3}$	$+\ \epsilon_{2,3}$	$+\ \epsilon_3$			
$D_4=\mu$	$+\cdots$	$+\ \epsilon_{0,4}$	$+\ \epsilon_{1,4}$	$+\ \epsilon_{2,4}$	$+\ \epsilon_{3,4}$	$+\ \epsilon_4$		
$D_5=\mu$	$+\cdots$	$+\ \epsilon_{0,5}$	$+\ \epsilon_{1,5}$	$+\ \epsilon_{2,5}$	$+\ \epsilon_{3,5}$	$+\ \epsilon_{4,5}$	$+\ \epsilon_5$	
$D_6=\mu$	$+\cdots$	$+\ \epsilon_{0,6}$	$+\ \epsilon_{1,6}$	$+\ \epsilon_{2,6}$	$+\ \epsilon_{3,6}$	$+\ \epsilon_{4,6}$	$+\ \epsilon_{5,6}$	$+\ \epsilon_6$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$ $\vdots$

incremental information $\epsilon_{t-i,\,t}$ could be in the form of the retailer's forecast and/or advance demand information (e.g., Heath and Jackson 1994, Gallego and Ozer 2001).

Given the demand model of (1), the forecast for demand D_t at the end of period $t-i$ $(i\ge 0)$ can be naturally defined as the conditional expectation given by

$$F_{t-i,t}=\mu+\sum_{j=i}^{\infty}\epsilon_{t-j,\,t}, \tag{2}$$

with F_{tt} being the actual demand D_t itself.

By some simple algebra, we can show $F_{t-i,\,t}=F_{t-i-1,\,t}+\epsilon_{t-i,\,t}$. Here, $\epsilon_{t-i,\,t}$ can be viewed as the forecast revision with regard to demand D_t made at the end of period $t-i$, and ϵ_{tt} is the final uncertainty resolved during period t after demand realization. We replace ϵ_{tt} with ϵ_t henceforth to suppress notation. A simple illustration of the MMFE process is given in Table 1.

The incremental information that the retailer obtains at the end of period t with regard to all future demands can be summarized in a forecast revision vector $\boldsymbol{\epsilon}_t=[\epsilon_t,\epsilon_{t,\,t+1},\epsilon_{t,\,t+2},\ldots]^T$. When the forecast horizon is finite, say H periods, the forecast revision vector can be effectively truncated by setting $\epsilon_{t,\,t+i}=0$ for all $i>H$. In addition to the demand model (1), we further assume that the forecast revision vector $\boldsymbol{\epsilon}_t$ is independent and identically distributed with a multivariate normal distribution $N(\mathbf{0},\boldsymbol{\Sigma})$, with the variance-covariance matrix given by $\boldsymbol{\Sigma}=E\{\boldsymbol{\epsilon}_t\boldsymbol{\epsilon}_t^T\}$.

The MMFE demand model described thus far generalizes many commonly used demand models in the literature, such as the i.i.d. normal demand model, the AR(1) model (Lee et al. 2000, Raghunathan 2001), the IMA(0,1,1) model (Graves 1999, Miyaoka and Hausman 2004), the general ARMA model (Gilbert 2005, Gaur et al. 2005), the linear state-space model (Aviv 2003), and the advance demand information model (Gallego and Ozer 2001).

The i.i.d. normal demand model can be obtained by simply setting $\epsilon_{t,\,t+i}=0$ for all $i>0$, which means no forecast revision information is available. In this case, the demand variability σ^2 is the same as the demand uncertainty.

If we let the forecast revision vector $\boldsymbol{\epsilon}_t=[\epsilon_t,\rho\epsilon_t,\rho^2\epsilon_t,\ldots]^T$ (with $0\le|\rho|<1$) for all t, then we have $D_t=\mu+\sum_{i=1}^{\infty}\epsilon_{t-i,\,t}+\epsilon_t=\mu+\rho\cdot\sum_{i=1}^{\infty}\epsilon_{t-i,\,t-1}+\epsilon_t=\mu+\rho\cdot(D_{t-1}-\mu)+\epsilon_t$, which means the demand follows the AR(1) process. Interpreting this the other way around, we see that the random shock ϵ_t in the AR(1) model reveals partial future forecast information in the form of $\epsilon_{t,\,t+i}=\rho^i\epsilon_t$ for $i>0$.

Next, if we let the forecast revision vector $\boldsymbol{\epsilon}_t=[\epsilon_t,\alpha\epsilon_t,\alpha\epsilon_t,\ldots]^T$ (with $0<\alpha\le 1$) for all t, then we have $D_t=\mu+\sum_{i=1}^{\infty}\epsilon_{t-i,\,t}+\epsilon_t=\mu+\alpha\cdot\sum_{i=1}^{\infty}\epsilon_{t-i}+\epsilon_t=D_{t-1}-(1-\alpha)\epsilon_{t-1}+\epsilon_t$, which means the demand follows the IMA(0,1,1) process. In this case, the random shock ϵ_t reveals partial future forecast information in the form of $\epsilon_{t,\,t+i}=\alpha\epsilon_t$ for $i>0$. Note that in this case, $\sigma^2=\sum_{i=0}^{\infty}\sigma_i^2=\sigma_0^2+\sum_{i=1}^{\infty}\alpha^2\sigma_0^2=\infty$. So only the conditional demand variability would be meaningful. The results derived from the MMFE model should be modified with the conditional expectation in this case.

Now suppose the forecast revision vector takes a more general form $\boldsymbol{\epsilon}_t=[\epsilon_t,\psi_1\epsilon_t,\psi_2\epsilon_t,\ldots]^T$, with ψ_i determined by

$$1+\psi_1B+\psi_2B^2+\cdots=\frac{\theta(B)}{\varphi(B)(1-B)^d}=\frac{1-\theta_1B-\theta_2B^2-\cdots-\theta_qB^q}{(1-\varphi_1B-\varphi_2B^2-\cdots-\varphi_pB^p)(1-B)^d},$$

where B is the time-series backshift operator. Then, by some algebra, we can show

$$\varphi(B)(1-B)^d(D_t-\mu)=\theta(B)\epsilon_t,$$

which is the general ARIMA(p,d,q) time-series model (Box et al. 1994). In this case, the random shock ϵ_t reveals partial future forecast information in the form of $\epsilon_{t,\,t+i}=\psi_i\epsilon_t$ for $i>0$. For the case when $d\ge 1$, from Box et al. (1994), $\sigma^2=\sum_{i=0}^{\infty}\sigma_i^2=\sigma_0^2+\sum_{i=1}^{\infty}\psi_i^2\sigma_0^2=\infty$. So the results derived from the MMFE model should be modified with the conditional expectation in this case.

Furthermore, suppose the forecast revision vector takes the following form:

$$\boldsymbol{\epsilon}_t=[RHV_t,RHFV_t,RHF^2V_t,\ldots]^T,$$

with R being a $1\times m$ vector, H an $m\times n$ matrix, F an $n\times n$ matrix, and V_t an n-dimensional vector. In addition, V_t is independent and identically distributed with a multivariate normal distribution. Then we can show by some algebra that the demand in a period is given by $D_t=\mu+R\Psi_t$, with the m-dimensional vector Ψ_t determined by the following system equations:

$$X_t=FX_{t-1}+V_t,$$

$$\Psi_t=HX_t,$$

where X_t is an n-dimensional state vector. So in this case, the MMFE demand model is the same as the linear state-space model (Aviv 2003). From the forecast revision vector, we can see that the partial future forecast information is revealed in the form of $\epsilon_{t,t+i} = RHF^iV_t$ for $i > 0$.

Finally, suppose the forecast revision vector takes the form of $\boldsymbol{\epsilon}_t = [\eta_0, \eta_1, \eta_2, \ldots]^T$, where η_i is mutually independent for all $i \geq 0$. If we view η_i as the incremental advance orders placed by customers (to be delivered in future periods), then this is essentially the same as the advance demand information model studied by Gallego and Ozer (2001). Note that the demand model in this case does not belong to either the ARIMA or the linear state-space model; so the MMFE model is a superset of the ARIMA and the linear state-space model.

3. A Generalized Order-Up-To Policy

In this section, we will first introduce a generalized order-up-to policy based on the MMFE demand model and then determine the supply chain costs under this policy.

3.1. Assumptions and Policy Structure

Let us first consider the retailer's ordering process. The replenishment lead time for the retailer is assumed to be constant at L periods. The timing of events is the following: (1) At the beginning of a period, the retailer places an order with the supplier. (2) Next, the goods ordered L periods ago arrive. (3) Finally, demand is realized, and the available inventory is used to meet the demand. Excess demand is backlogged. Unit holding cost h and penalty cost p are assessed and charged to the retailer at the end of each period.

Following the derivation of the previous section, we observe that prior to the beginning of a period t, all the information that the retailer possesses about the future demands is given in the form of a series of forecast revision vectors up until period $t-1$, i.e., $\{\boldsymbol{\epsilon}_{-\infty}, \ldots, \boldsymbol{\epsilon}_{t-1}\}$.

If the retailer optimizes its own cost, we can show that under some standard assumptions, the optimal order-up-to level is an affine time-invariant combination of historical forecast revision vectors (see Proposition 1 in §3.2). This affine updating is essentially being used in most commercial software (such as Evant and Manugistics). Because our goal is to propose an order-smoothing policy that the retailer can use to minimize the total supply chain costs, it is natural to generalize the order-up-to policy by preserving the elegant structure of affinity and stationarity. Specifically, we will consider a class of order-up-to policies with the following properties: (a) stationarity—the mapping from the forecast revision vectors to the order-up-to level is time-invariant; (b) affinity—the mapping from the forecast revision vectors to the order-up-to level is affine.

Given these two properties, the order-up-to level can be defined in the following format:

$$S_t = m + \sum_{i=1}^{\infty} \mathbf{w}_i^T \boldsymbol{\epsilon}_{t-i}, \tag{3}$$

where m is a constant and $\mathbf{w}_i$ is a weight vector defined as $\mathbf{w}_i = [w_{i1}, w_{i2}, w_{i3}, \ldots]^T$ for $i > 0$ (no restrictions are imposed on the value of $\mathbf{w}_i$). For notational convenience, let us define $\mathbf{w}_0 = \mathbf{0}$.

The generalized order-up-to policy described above can be viewed as a class of order-smoothing policy in the sense that the weight vector $\mathbf{w}_i$ plays a role in smoothing the forecast revision information to produce a desirable inventory order-up-to level. Graves et al. (1998) and Aviv (2007) derived smoothing policies by minimizing certain production-smoothing costs (such as the cost of production variability and adherence-to-plans). Their production-smoothing policies also bear a linear (or affine) time-invariant structure, which in a sense corroborates our two assumptions on the generalized order-up-to policy.

Under this order-up-to policy, the retailer's order quantity at the beginning of period t is

$$O_t = S_t - (S_{t-1} - D_{t-1}) = \mu + \sum_{i=1}^{\infty} (\mathbf{w}_i - \mathbf{w}_{i-1} + \mathbf{e}_i)^T \boldsymbol{\epsilon}_{t-i}, \tag{4}$$

where $\mathbf{e}_i$ is the unit vector with the ith element equal to one. We further define the following notation: $\mathbf{e}_i^j = \sum_{k=i}^{j} \mathbf{e}_k$. This notation will be used throughout the paper.

We note that under this order policy, the order quantity O_t is not guaranteed to be positive because of the normality assumption of the forecast revision vector $\boldsymbol{\epsilon}_t$. When O_t is restricted to be positive, exact analysis of the system becomes difficult. Approximate approaches to this problem can be found in papers by Iida and Zipkin (2006) and Lu et al. (2006). Alternatively, to ensure $O_t \geq 0$ for all t, we could impose stronger assumptions on the distribution of the forecast revision vector. For example, using a similar argument of Johnson and Thompson (1975), we can assume that $\epsilon_{t,t+i}$ only takes value in a certain range $(-a_i, a_i)$ and that μ is sufficiently large such that $\mu > \sum_{i=1}^{\infty} |\mathbf{w}_i - \mathbf{w}_{i-1} + \mathbf{e}_i|^T \mathbf{a}$, where $\mathbf{a} = [a_0, a_1, \ldots]^T$. Then from (4), we have $O_t \geq 0$ with probability one. But imposing this kind of restrictive assumption would introduce significant notational and analytical complexity. So for the sake of ease of exposition and tractability, we will henceforth allow for negative O_t. In other words, we effectively assume that the retailer can freely return excess inventory to the supplier (see Kahn 1987, Lee et al. 1997).

The free-return assumption becomes quite innocuous when the demand mean is sufficiently large (e.g., $\mu \gg \sqrt{\text{var}(O_t)}$) because the chance of a negative order quantity will become negligible. For example, if the demand follows an AR(1) model with parameter $|\rho| < 1$, by some algebra, one can show that for $\rho \geq 0$, $\sqrt{\text{var}(O_t)} \leq \sqrt{\mathbf{1}^T \boldsymbol{\Sigma} \mathbf{1}} = \sqrt{(1+\rho)/(1-\rho)} \cdot \sigma$. Hence, as long as $\mu \gg \sqrt{(1+\rho)/(1-\rho)} \cdot \sigma$, we can safely ignore the occurrence of negative order quantity. For a reasonably small ρ value, this is just slightly more restrictive than the condition $\mu \gg \sigma$ as required to avoid a negative demand under the normality assumption. Similar justifications for the free-return assumption are also provided by Lee et al. (2000) and Aviv (2003). In particular, Aviv (2007) found that the rate of occurrence of a negative production quantity is less than 1% across all simulation instances under a production-smoothing model with quadratic cost function.

3.2. The Retailer's Cost

Let us define the retailer's single-period cost function as $c(x) = h \cdot \max(x, 0) + p \cdot \max(-x, 0)$. Assume that the retailer's order quantity is never shorted by the supplier. Then under the generalized order-up-to policy, the retailer's long-run average cost is given by

$$
\begin{aligned}
&C(m, \mathbf{w}_1, \mathbf{w}_2, \ldots) \\
&= E\left\{c\left(S_t - \sum_{i=0}^{L} D_{t+i}\right)\right\} \\
&= E\left\{c\left(m + \sum_{i=1}^{\infty} \mathbf{w}_i^T \boldsymbol{\epsilon}_{t-i} - (L+1)\mu - \sum_{i=0}^{L}\sum_{j=0}^{\infty} \boldsymbol{\epsilon}_{t+i-j,\, t+i}\right)\right\} \\
&= E\left\{c\left(m - (L+1)\mu + \sum_{i=1}^{\infty} (\mathbf{w}_i - \mathbf{e}_{i+1}^{i+L+1})^T \boldsymbol{\epsilon}_{t-i} \right.\right. \\
&\qquad \left.\left. - \sum_{i=0}^{L}\sum_{j=0}^{i} \boldsymbol{\epsilon}_{t+i-j,\, t+i}\right)\right\}.
\end{aligned}
$$

Let us fix the weight vector $\mathbf{w}_i$ and find the optimal m to minimize $C(m, \mathbf{w}_1, \mathbf{w}_2, \ldots)$. From the above expression, it is clear that there are two mutually independent sources of variability: $\sum_{i=1}^{\infty} (\mathbf{w}_i - \mathbf{e}_{i+1}^{i+L+1})^T \boldsymbol{\epsilon}_{t-i}$ and $\sum_{i=0}^{L}\sum_{j=0}^{i} \boldsymbol{\epsilon}_{t+i-j,\, t+i}$. The first source of variability is a function of the weight vector $\mathbf{w}_i$. Let $\Delta_w(\mathbf{w}_1, \mathbf{w}_2, \ldots)$ denote its variance; i.e.,

$$
\begin{aligned}
\Delta_w(\mathbf{w}_1, \mathbf{w}_2, \ldots) &= \text{var}\left(\sum_{i=1}^{\infty} (\mathbf{w}_i - \mathbf{e}_{i+1}^{i+L+1})^T \boldsymbol{\epsilon}_{t-i}\right) \\
&= \sum_{i=1}^{\infty} (\mathbf{w}_i - \mathbf{e}_{i+1}^{i+L+1})^T \boldsymbol{\Sigma} (\mathbf{w}_i - \mathbf{e}_{i+1}^{i+L+1}).
\end{aligned}
$$

The second source of variability is the exposure-period uncertainty that is resolved by the forecast information. Let Δ_d denote its variance; i.e.,

$$
\Delta_d = \text{var}\left(\sum_{i=0}^{L}\sum_{j=0}^{i} \epsilon_{t+i-j,\, t+i}\right) = \sum_{i=1}^{L+1} \mathbf{e}_1^{i\,T} \boldsymbol{\Sigma} \mathbf{e}_1^i. \tag{5}
$$

This is a classic newsvendor problem under normal distribution, so it is straightforward to show that $m^*(\mathbf{w}_1, \mathbf{w}_2, \ldots) = (L+1)\mu + z\sqrt{\Delta_w(\mathbf{w}_1, \mathbf{w}_2, \ldots) + \Delta_d}$, where $z = \Phi^{-1}(p/(h+p))$ ($\Phi(\cdot)$ is the standard normal cumulative distribution function). The resulting optimal cost can be shown as

$$
\begin{aligned}
&C(m^*, \mathbf{w}_1, \mathbf{w}_2, \ldots) \\
&\qquad = (h+p)\phi(z)\sqrt{\Delta_w(\mathbf{w}_1, \mathbf{w}_2, \ldots) + \Delta_d},
\end{aligned} \tag{6}
$$

where $\phi(\cdot)$ is the standard normal density function (Porteus 2002). From the expression, it is clear that the retailer's cost is minimized when $\Delta_w(\mathbf{w}_1, \mathbf{w}_2, \ldots) = 0$.

Proposition 1. *Assume that the retailer can freely return unsold inventory and the supplier has ample supply. If the retailer optimizes its own cost, the optimal order-up-to level S_t is*

$$
\begin{aligned}
S_t &= \sum_{i=0}^{L} F_{t-1,\, t+i} + z\sqrt{\Delta_d} \\
&= \sum_{i=1}^{\infty} (\mathbf{e}_{i+1}^{i+L+1})^T \boldsymbol{\epsilon}_{t-i} + (L+1)\mu + z\sqrt{\Delta_d},
\end{aligned}
$$

which is an affine time-invariant function of forecast revision vectors $\{\boldsymbol{\epsilon}_{-\infty}, \ldots, \boldsymbol{\epsilon}_{t-1}\}$ with the weight vector $\mathbf{w}_i^* = \mathbf{e}_{i+1}^{i+L+1}$ for $i > 0$ and the constant $m^* = (L+1)\mu + z\sqrt{\Delta_d}$. Moreover, the order variability amplification ratio in this case is given by

$$
\begin{aligned}
\frac{\text{var}(O_t)}{\text{var}(D_{t-1})} &= 1 + \frac{\text{var}(\sum_{i=0}^{L+1} \boldsymbol{\epsilon}_{t,\, t+i}) - \sum_{i=0}^{L+1} \text{var}(\boldsymbol{\epsilon}_{t,\, t+i})}{\sigma^2} \\
&= 1 + \frac{(\mathbf{e}_1^{L+2})^T \boldsymbol{\Sigma} \mathbf{e}_1^{L+2} - \sum_{i=1}^{L+2} (\mathbf{e}_i)^T \boldsymbol{\Sigma} \mathbf{e}_i}{\sigma^2}.
\end{aligned}
$$

The above proposition shows that if the retailer only optimizes its own cost, the optimal order-up-to level is simply the demand forecast over the exposure period plus the safety stock quantity required to cover the demand uncertainty Δ_d. This order-up-to level is also an affine time-invariant function of historical forecast revision vectors, which supports our two assumptions on the generalized order-up-to policy.

Furthermore, under the general MMFE demand model, we obtain a general formula for the order variability amplification ratio. We note that this amplification ratio is based on the unconditional variance of demand and order, which is commonly used in the literature to measure the bullwhip effect. If we substitute the forecast revision vector with the specific form for the AR(1) model (see §2), i.e., $\boldsymbol{\epsilon}_t = [\epsilon_t, \rho\epsilon_t, \rho^2\epsilon_t, \ldots]^T$, then we can easily recover the original bullwhip effect result derived by Lee et al. (1997) as follows:

$$
\frac{\text{var}(O_t)}{\text{var}(D_{t-1})} = 1 + \frac{2\rho(1-\rho^{L+1})(1-\rho^{L+2})}{1-\rho}.
$$

Similarly, for all other time-series demand models, the bullwhip effect results can be unified under the general formula given in Proposition 1.

From the AR(1) formula, the bullwhip effect can be attributed to either long lead time or high autocorrelation (Lee et al. 1997). However, under the general formula given in Proposition 1, we see that increasing lead time alone does not necessarily ensure the increase of the bullwhip effect. It is actually the forecast correlation over the exposure period as a whole that determines the magnitude of the bullwhip effect. In this sense, we obtain a more precise insight about the cause of the bullwhip effect under the generalized demand model.

3.3. The Supplier's Cost

To evaluate the supplier's cost, we need to consider the supplier's order fulfillment process. Let us assume that whenever the supplier runs out of inventory, it will draw from a secondary source to cover the shortfall amount. Therefore, the retailer will always receive a full order from the supplier, but part of it may come from the secondary source. Under this assumption, which was also made in previous studies (Gavirneni et al. 1999, Lee et al. 2000, Balakrishnan et al. 2004), the retailer's long-run average cost is the same as that of the ample supply case discussed in the previous section. Naturally, the supplier would incur a unit cost for using the secondary source. Let p' denote such unit penalty cost. In addition, let h' denote the unit inventory holding cost for the supplier.

The production lead time for the supplier is assumed to be constant at l periods. The timing of the supplier's order-fulfilling process is given as follows. (1) Before the beginning of a period, the goods ordered l period ago arrive. (2) At the beginning of a period, the supplier receives and ships the required order quantity to the retailer. If the supplier does not have enough stock to fill this order, the shortfall is filled by a secondary source. (3) The supplier reviews its inventory level and places an order with its external supplier. We assume that the supplier's external supplier and the secondary source always have ample supply.

When the supplier's production lead time l is longer than one period, exact analysis of the supplier's exposure-period demand distribution becomes cumbersome, as it involves accounting of the shortage amounts that are fulfilled by the secondary resource in each period. In this case, a simplification approach is to assume that the shortfall inventory is "borrowed" from the secondary source, which needs to be returned to the secondary source after usage (see Lee et al. 2000). This simplifying assumption effectively turns the supplier's inventory problem into a much easier full-backlog problem. On the other hand, when the supplier's lead time is one period, the analysis of the supplier's exposure-period demand distribution becomes straightforward, and there is no need to impose the simplifying "borrowed" inventory assumption.

In what follows, for ease of exposition, we will assume the supplier's production lead time $l = 1$ (see Bourland et al. 1996 and Gavirneni et al. 1999 for similar treatments). Extending our results to the case of $l > 1$ is straightforward with the aid of the "borrowed" inventory assumption. This extension will be discussed in detail in §6.

Let S'_t denote the supplier's order-up-to level. Analogous to (3), we define the following:

$$S'_t = m' + \sum_{i=1}^{\infty} \mathbf{w}'^T_i \boldsymbol{\epsilon}_{t-i}, \tag{7}$$

where m' is a constant and $\mathbf{w}'_i$ $(i > 0)$ is the supplier's weight vector. We also define $\mathbf{w}'_0 = \mathbf{0}$ for notational convenience. When the supplier does not know the forecast revision vector $\boldsymbol{\epsilon}_t$, we can set $\mathbf{w}'_i = \mathbf{0}$ for $i > 0$. In this case, the order-up-to level reduces to a static one.

Define the supplier's single-period cost function as $c'(x) = h' \cdot \max(x, 0) + p' \cdot \max(-x, 0)$. Recall the retailer's order quantity as given in (4). The supplier's long-run average cost is thus given by

$$\begin{aligned}
&C'(m', \mathbf{w}'_1, \mathbf{w}'_2, \ldots, \mathbf{w}_1, \mathbf{w}_2, \ldots)\\
&= E\{c'(S'_t - O_{t+1})\}\\
&= E\left\{c'\left(m' + \sum_{i=1}^{\infty} \mathbf{w}'^T_i \boldsymbol{\epsilon}_{t-i} - \mu - \sum_{i=1}^{\infty} (\mathbf{w}_i - \mathbf{w}_{i-1} + \mathbf{e}_i)^T \boldsymbol{\epsilon}_{t+1-i}\right)\right\}\\
&= E\left\{c'\left(m' - \mu + \sum_{i=1}^{\infty} (\mathbf{w}'_i - \mathbf{w}_{i+1} + \mathbf{w}_i - \mathbf{e}_{i+1})^T \boldsymbol{\epsilon}_{t-i}\right.\right.\\
&\qquad\qquad \left.\left. - (\mathbf{w}_1 + \mathbf{e}_1)^T \boldsymbol{\epsilon}_t\right)\right\}.
\end{aligned} \tag{8}$$

From the above expression, we see that there are two mutually independent sources of variability: $\sum_{i=1}^{\infty} (\mathbf{w}'_i - \mathbf{w}_{i+1} + \mathbf{w}_i - \mathbf{e}_{i+1})^T \boldsymbol{\epsilon}_{t-i}$ and $(\mathbf{w}_1 + \mathbf{e}_1)^T \boldsymbol{\epsilon}_t$. The first source of variability is a function of $\mathbf{w}'_i$, which can be potentially eliminated if the supplier knows the retailer's forecast revision vector $\boldsymbol{\epsilon}_{t-i}$ and weight vector $\mathbf{w}_i$. When there is no information sharing, we have $\mathbf{w}'_i = \mathbf{0}$ for $i > 0$. Let $\Delta_f(\mathbf{w}_1, \mathbf{w}_2, \ldots)$ denote the variance in this case; i.e.,

$$\begin{aligned}
&\Delta_f(\mathbf{w}_1, \mathbf{w}_2, \ldots)\\
&\quad = \operatorname{var}\left(\sum_{i=1}^{\infty} (\mathbf{w}_{i+1} - \mathbf{w}_i + \mathbf{e}_{i+1})^T \boldsymbol{\epsilon}_{t-i}\right)\\
&\quad = \sum_{i=1}^{\infty} (\mathbf{w}_{i+1} - \mathbf{w}_i + \mathbf{e}_{i+1})^T \boldsymbol{\Sigma} (\mathbf{w}_{i+1} - \mathbf{w}_i + \mathbf{e}_{i+1}).
\end{aligned} \tag{9}$$

The second source of variability is the order uncertainty that cannot be resolved by information sharing. Let $\Delta_o(\mathbf{w}_1)$ denote its variance; i.e.,

$$\Delta_o(\mathbf{w}_1)=\mathrm{var}((\mathbf{w}_1+\mathbf{e}_1)^T\boldsymbol{\epsilon}_t)=(\mathbf{w}_1+\mathbf{e}_1)^T\boldsymbol{\Sigma}(\mathbf{w}_1+\mathbf{e}_1). \quad (10)$$

By analyzing the supplier's cost function, we can quantify the value of information sharing.

4. Value of Information Sharing

A general consensus in the literature is that information sharing between a retailer and supplier would benefit the supplier. The magnitude of such benefit, however, differs considerably depending on various modeling assumptions. For example, Lee et al. (2000) showed that under the AR(1) demand model, sharing POS information would provide significant value to the supplier. In contrast, under the same model, Raghunathan (2001) showed that if the supplier utilizes all the order history information from the retailer, then sharing POS information does not provide much value to the supplier. In the following, with the aid of our generalized demand model and order-up-to policy, we will provide a generic quantification of the value of information sharing and then reconcile the results of these studies.

Let us consider two scenarios. First, if there is no information sharing, it is equivalent to have $\mathbf{w}'_i=\mathbf{0}$ for $i>0$. In this case, from expression (8), the supplier's cost is given by

$$C'_N(m',\mathbf{w}_1,\mathbf{w}_2,\ldots)=E\left\{c'\left(m'-\mu-\sum_{i=1}^{\infty}(\mathbf{w}_{i+1}-\mathbf{w}_i+\mathbf{e}_{i+1})^T\boldsymbol{\epsilon}_{t-i} - (\mathbf{w}_1+\mathbf{e}_1)^T\boldsymbol{\epsilon}_t\right)\right\}.$$

The optimal m' for this newsvendor cost function is $m'^*_N=\mu+z'\sqrt{\Delta_f(\mathbf{w}_1,\mathbf{w}_2,\ldots)+\Delta_o(\mathbf{w}_1)}$, where $z'=\Phi^{-1}(p'/(h'+p'))$ and $\Delta_f(\mathbf{w}_1,\mathbf{w}_2,\ldots)$ and $\Delta_o(\mathbf{w}_1)$ are defined in (9) and (10), respectively. The resulting optimal cost for this no-information-sharing case is

$$C'_N(m'^*_N,\mathbf{w}_1,\mathbf{w}_2,\ldots) = (h'+p')\phi(z')\sqrt{\Delta_f(\mathbf{w}_1,\mathbf{w}_2,\ldots)+\Delta_o(\mathbf{w}_1)}.$$

Second, if the retailer shares the forecast revision information $\boldsymbol{\epsilon}_t$ as well as the order weight vector $\mathbf{w}_i$ with the supplier, then the supplier can reduce its cost by setting $\mathbf{w}'_i=\mathbf{w}_{i+1}-\mathbf{w}_i+\mathbf{e}_{i+1}$ for $i>0$ in expression (8). The supplier's cost in this information sharing case is then given by

$$C'_I(m',\mathbf{w}_1)=E\{c'(m'-\mu-(\mathbf{w}_1+\mathbf{e}_1)^T\boldsymbol{\epsilon}_t)\}. \quad (11)$$

The optimal m' for this newsvendor cost function is $m'^*_I=\mu+z'\sqrt{\Delta_o(\mathbf{w}_1)}$, and the resulting optimal cost with information sharing is

$$C'_I(m'^*_I,\mathbf{w}_1)=(h'+p')\phi(z')\sqrt{\Delta_o(\mathbf{w}_1)}. \quad (12)$$

Proposition 2. *If the retailer shares with the supplier the forecast revision information $\boldsymbol{\epsilon}_t$ and the order weight vector $\mathbf{w}_i$ ($i>0$), then the net cost savings for the supplier are given by*

$$(h'+p')\phi(z')\left(\sqrt{\Delta_f(\mathbf{w}_1,\mathbf{w}_2,\ldots)+\Delta_o(\mathbf{w}_1)}-\sqrt{\Delta_o(\mathbf{w}_1)}\right), \quad (13)$$

where $z'=\Phi^{-1}(p'/(h'+p'))$, $\Delta_f(\mathbf{w}_1,\mathbf{w}_2,\ldots)$, and $\Delta_o(\mathbf{w}_1)$ are defined in (9) *and* (10), *respectively.*

Recall that in §2, we show that the forecast revision vector for the AR(1) model is $\boldsymbol{\epsilon}_t=[\epsilon_t,\rho\epsilon_t,\rho^2\epsilon_t,\ldots]^T$. According to Lee et al. (2000), sharing the POS information in period t is the same as sharing ϵ_t. Therefore, assuming that the supplier has the knowledge of the model parameter ρ, sharing POS information ϵ_t is equivalent to sharing the forecast revision vector $\boldsymbol{\epsilon}_t$. Hence, it is not surprising that Lee et al. (2000) concluded that sharing POS information provides significant value to the supplier.

Suppose that the retailer optimizes its own cost (same as in Lee et al. 2000). From Proposition 1, the optimal weight vector in this case is $\mathbf{w}_i=\mathbf{e}_{i+1}^{i+L+1}$. Therefore, the cost savings for the supplier under the AR(1) model are, according to Proposition 2,

$$\begin{aligned}(h'+p')\phi(z')&\left(\sqrt{\Delta_f(\mathbf{e}_2^{L+2},\mathbf{e}_3^{L+3},\ldots)+\Delta_o(\mathbf{e}_2^{L+2})}-\sqrt{\Delta_o(\mathbf{e}_2^{L+2})}\right)\\ &=(h'+p')\phi(z')\left(\sqrt{\sum_{i=1}^{\infty}(\mathbf{e}_{i+L+2})^T\boldsymbol{\Sigma}\mathbf{e}_{i+L+2}+(\mathbf{e}_1^{L+2})^T\boldsymbol{\Sigma}\mathbf{e}_1^{L+2}} - \sqrt{(\mathbf{e}_1^{L+2})^T\boldsymbol{\Sigma}\mathbf{e}_1^{L+2}}\right) \quad (14)\\ &=(h'+p')\phi(z')\sigma_0\left(\sqrt{\frac{\rho^{2(L+2)}}{1-\rho^2}+\left(\frac{1-\rho^{L+2}}{1-\rho}\right)^2}-\frac{1-\rho^{L+2}}{1-\rho}\right), \quad (15)\end{aligned}$$

where the last step follows from the fact that

$$\boldsymbol{\Sigma}=\sigma_0^2\cdot\begin{pmatrix}1&\rho&\rho^2&\cdots\\ \rho&\rho^2&\rho^3&\cdots\\ \rho^2&\rho^3&\rho^4&\cdots\\ \vdots&\vdots&\vdots&\ddots\end{pmatrix}.$$

The result derived in (15) is the same as shown by Lee et al. (2000). Analogously, for a general ARMA(p,q) model, we know from §2 that the forecast revision vector is given by $\boldsymbol{\epsilon}_t=[\epsilon_t,\psi_1\epsilon_t,\psi_2\epsilon_t,\ldots]^T$, with ψ_i determined by

$$1+\psi_1B+\psi_2B^2+\cdots=\frac{\theta(B)}{\varphi(B)} =\frac{1-\theta_1B-\theta_2B^2-\cdots-\theta_qB^q}{(1-\varphi_1B-\varphi_2B^2-\cdots-\varphi_pB^p)},$$

where B is the time-series backshift operator. Therefore, in this case, we have

$$\mathbf{\Sigma} = \sigma_0^2 \cdot \begin{pmatrix} 1 & \psi_1 & \psi_2 & \cdots \\ \psi_1 & \psi_1^2 & \psi_1\psi_2 & \cdots \\ \psi_2 & \psi_2\psi_1 & \psi_2^2 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}.$$

Substituting the above $\mathbf{\Sigma}$ into (14), we obtain the cost savings for the supplier under the general ARMA(p, q) model as

$$(h'+p')\phi(z')\sigma_0\left(\sqrt{\sum_{i=L+2}^{\infty}\psi_i^2 + \left(1+\sum_{i=1}^{L+1}\psi_i\right)^2} - \sqrt{\left(1+\sum_{i=1}^{L+1}\psi_i\right)^2}\right).$$

This result generalizes the AR(1) case and the ARMA(1, 1) case studied by Gaur et al. (2005).

On the other hand, if the supplier knows that the retailer's forecast revision vector is of the form $\boldsymbol{\epsilon}_t = [\epsilon_t, \rho\epsilon_t, \rho^2\epsilon_t, \ldots]^T$ and the weight vector is $\mathbf{w}_i = \mathbf{e}_{i+1}^{i+L+1}$, then by (4) the supplier can deduce that

$$O_t = \sum_{i=1}^{\infty}\rho^{L+i+1}\epsilon_{t-1-i} + \frac{1-\rho^{L+2}}{1-\rho}\epsilon_{t-1}$$

for all t. With some algebra, it can be shown

$$\epsilon_{t-1} = \frac{1-\rho}{1-\rho^{L+2}}O_t - \frac{\rho-\rho^2}{1-\rho^{L+2}}O_{t-1} + \frac{\rho-\rho^{L+2}}{1-\rho^{L+2}}\epsilon_{t-2}.$$

By repeating this deduction for ϵ_{t-2} and so on, ϵ_{t-1} can be asymptotically approximated by a linear combination of historical orders $\{O_t, O_{t-1}, \ldots\}$ as long as $|(\rho-\rho^{L+2})/(1-\rho^{L+2})| < 1$ (or $-0.5 < \rho < 1$). This is essentially the same argument used in Corollary 1 of Gaur et al. (2005). Therefore, the forecast revision vector can be inferred by using just the historical order information. The resulting cost savings for the supplier would approach to the maximum possible level given in (15). This explains the result obtained by Raghunathan (2001): sharing POS information under the AR(1) model provides negligible value because the supplier can infer all the information from the historical orders. By analogous arguments, for a general ARMA(p, q) model, we can explain the result obtained by Gaur et al. (2005) that if the order process (4) is invertible, then the forecast revision information is inferable from the historical orders. Such inference, however, would not be possible if the supplier does not have the knowledge of the retailer's demand model and order policy. So one can still argue that some form of information sharing between the retailer and the supplier is required to make this work.

From the above analysis, we have observed that for the supplier to fully realize the benefit of information such as POS or raw order data, the supplier needs to have the knowledge of the retailer's demand model and order policy. But in reality, such knowledge is usually not available to the supplier. Even if it is available, the supplier would still need to be sophisticated enough to use this knowledge to realize the value of information. So in this sense, the commonly used "omniscient supplier" assumption in the information sharing literature (e.g., Lee et al. 2000, Raghunathan 2001, Aviv 2003, Gaur et al. 2005) is an aggressive one to use in practice. Removing this assumption would require the retailer to share more relevant and easy-to-use data with the supplier, so that the supplier can fully benefit from the information.

Specifically, to eliminate the need for the supplier to process and interpret the retailer's demand model and order policy, we propose an effective way of information sharing—the retailer should share its *projected future orders* with the supplier. This can be achieved by having the retailer share with the supplier the projections of future orders, where those projections can be revised over time.

Let $o_{t-i,t}$ denote the advance order revision made i periods before the order quantity O_t is finalized in period t (think $o_{t-i,t}$ to the supplier as the forecast revision $\epsilon_{t-i,t}$ to the retailer). We also replace o_{tt} with o_t to suppress notation. If the retailer places a base order of μ for each of the future periods and then revises its future order projections according to the following:

$$o_{t-i,t} = (\mathbf{w}_{i+1} - \mathbf{w}_i + \mathbf{e}_{i+1})^T\boldsymbol{\epsilon}_{t-1-i}, \tag{16}$$

then the final order quantity for period t is

$$O_t = \mu + \sum_{i=0}^{\infty}o_{t-i,t} = \mu + \sum_{i=1}^{\infty}(\mathbf{w}_i - \mathbf{w}_{i-1} + \mathbf{e}_i)^T\boldsymbol{\epsilon}_{t-i},$$

which is the same as that given in (4).

In this case, the supplier would receive an order revision vector $\mathbf{o}_t = [o_t, o_{t,t+1}, o_{t,t+2}, \ldots]^T$ in each period. The MMFE demand process faced by the retailer is effectively transformed into an MMFE order process for the supplier. Similar characteristics of the order process were also observed by Aviv (2003) and Gilbert (2005) for the linear state-space model and the ARIMA model, respectively.

Given all the historical order revision vectors $\{\mathbf{o}_{-\infty}, \ldots, \mathbf{o}_t\}$, we can redefine the supplier's order-up-to level S_t' as

$$S_t' = m' + \sum_{i=0}^{\infty}\mathbf{w'}_{i+1}^{T}\mathbf{o}_{t-i}.$$

Note that here the summation index starts at $i = 0$ rather than $i = 1$ as shown in (7) because $\mathbf{o}_t$ is based on $\boldsymbol{\epsilon}_{t-1}$ and the retailer places projected orders before the supplier ordering in a period.

The supplier's long-run average cost is thus given by

$$\begin{aligned} C'(m', \mathbf{w}_1', \mathbf{w}_2', \ldots) &= E\{c'(S_t' - O_{t+1})\} \\ &= E\left\{c'\left(m' + \sum_{i=0}^{\infty} \mathbf{w}'^T_{i+1} \mathbf{o}_{t-i} - \mu - \sum_{i=0}^{\infty} o_{t+1-i,\,t+1}\right)\right\} \\ &= E\left\{c'\left(m' - \mu + \sum_{i=0}^{\infty} (\mathbf{w}'_{i+1} - \mathbf{e}_{i+2})^T \mathbf{o}_{t-i} - o_{t+1}\right)\right\}. \end{aligned} \quad (17)$$

By setting $\mathbf{w}_i' = \mathbf{e}_{i+1}$ for all $i > 0$ in (17), we obtain $C'(m', \mathbf{w}_1', \mathbf{w}_2', \ldots) = E\{c'(m' - \mu - o_{t+1})\}$. The stochastic characterization of the order-revision quantity o_{t+1} can be determined based on the historical order-revision information $\{o_t, o_{t-1}, \ldots\}$, which is fairly straightforward to do. Note that from (16), we know that $o_{t+1} = (\mathbf{w}_1 + \mathbf{e}_1)^T \boldsymbol{\epsilon}_t$. Suppose that the stochastic characterization of o_{t+1} can be accurately determined; then the supplier's problem under the sharing of the order projections becomes the same as the information-sharing case $C_I'(m', \mathbf{w}_1)$ given in (11). Therefore, the resulting optimal cost for the supplier is the same as given in (12). The following proposition summarizes the above discussion:

Proposition 3. *If the retailer shares with the supplier the projected future order revision information* $\mathbf{o}_t = [o_t, o_{t,t+1}, o_{t,t+2}, \ldots]^T$ *as specified in the following for* $i \geq 0$,

$$o_{t,t+i} = (\mathbf{w}_{i+1} - \mathbf{w}_i + \mathbf{e}_{i+1})^T \boldsymbol{\epsilon}_{t-1},$$

and if the supplier uses the historical order-revision information to determine the variance-covariance structure of the order-revision process, then the net cost savings for the supplier will be the same as given in (13), *i.e., the same as if the retailer were sharing the forecast revision information* $\boldsymbol{\epsilon}_t$ *and the order weight vector* $\mathbf{w}_i$ $(i > 0)$.

In an empirical study, Terwiesch et al. (2005) reported that sharing projected orders is a common practice in a semiconductor equipment supply chain. Although this practice is reported to help better allocate scarce capacity, Proposition 3 shows that sharing projected future orders can also help the supplier achieve the full value of information. Therefore, instead of pushing for customer demand (POS) data or demand forecast sharing, an industry should also embrace the sharing of order projections.

5. Order Variability Control

Let us now assume that information sharing between the retailer and the supplier is achieved by the sharing of projected orders, as described in §4. The next question is how to determine the optimal weight vector $\mathbf{w}_i$ to minimize the total supply chain costs. Intuitively, with the supplier's lead time being one period, if the retailer told the supplier what it would order the next period and not revise it afterward (i.e., absorb the cost of uncertainty), the supplier's cost would be minimized. So, to balance the costs between the retailer and the supplier, the intuition is to have the retailer revise the order quantity only partially. Below we show this is indeed the case.

Let us combine (6) and (12) to obtain the total supply chain costs as

$$\begin{aligned} C(m^*, \mathbf{w}_1, \mathbf{w}_2, \ldots) + C_I'(m_I'^*, \mathbf{w}_1) &= (h+p)\phi(z)\sqrt{\Delta_w(\mathbf{w}_1, \mathbf{w}_2, \ldots) + \Delta_d} \\ &\quad + (h'+p')\phi(z')\sqrt{\Delta_o(\mathbf{w}_1)}. \end{aligned} \quad (18)$$

Define the relative cost ratio between the retailer and the supplier as

$$k = \frac{(h+p)\phi(z)}{(h'+p')\phi(z')}. \quad (19)$$

Note that k can be written as

$$k = \frac{h}{h'} \cdot \frac{(1+p/h)\phi(\Phi^{-1}(p/h/(1+p/h)))}{(1+p'/h')\phi(\Phi^{-1}(p'/h'/(1+p'/h')))}. \quad (20)$$

Figure 1 gives a visualization of the relationship between $(h+p)\phi(z)$ (with $z = \Phi^{-1}(p/(h+p))$) and the cost parameter ratio p/h. Inventory cost is normalized at $h = 1$ and the stockout penalty cost p is varied from 1 to 50, which represents a critical fractile from 50% to 98%. From the figure, we see that as the penalty cost p increases, the cost term $(h+p)\phi(z)$ is tripled from 0.8 to slightly greater than 2.4. Hence, according to (20), if we fix the ratios h/h' and p'/h',

Figure 1 **Illustration of $(h+p)\phi(z)$ (with $h = 1$)**

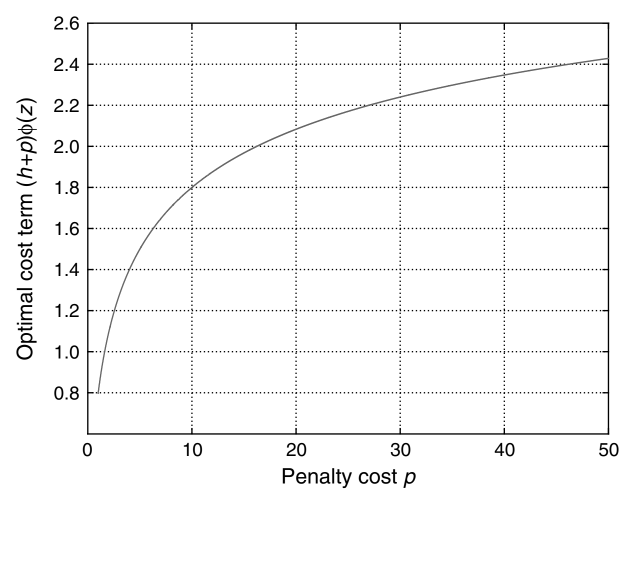

then k is increasing in the retailer's penalty-cost-to-holding-cost ratio p/h. Similarly, if we fix the ratios h/h' and p/h, then k is decreasing in the supplier's penalty-cost-to-holding-cost ratio p'/h'.

Also define the relative uncertainty ratio between the retailer and the supplier as

$$v = \frac{\Delta_d}{\Delta_o(\mathbf{e}_2^{L+2})}, \tag{21}$$

where Δ_d, defined in (5), is the total exposure-period demand uncertainty that is yet resolved by forecast information, and $\Delta_o(\mathbf{e}_2^{L+2})$ is the order uncertainty faced by the supplier when the retailer optimizes its own cost. The optimal weight vector is given in the following proposition:

Proposition 4. *The optimal weight vector* $\mathbf{w}_i^*$ *that minimizes the total supply chain costs* (18) *is*

$$\mathbf{w}_i^* = \begin{cases} -\gamma \mathbf{e}_1 + (1-\gamma)\mathbf{e}_2^{L+2} & i = 1, \\ \mathbf{e}_{i+1}^{i+L+1} & i \geq 2, \end{cases}$$

where γ *is given by*

$$\gamma = \begin{cases} \sqrt{\dfrac{v}{k^2-1}} & \text{if } k > \sqrt{1+v}, \\ 1 & \text{if } k \leq \sqrt{1+v}, \end{cases}$$

with k and v defined in (19) *and* (21), *respectively.*

Substituting the optimal weight vector $\mathbf{w}_i^*$ into the retailer's future order revision quantity $o_{t,\,t+i}$, as specified in Proposition 3, we derive that

$$o_{t,\,t+i}^* = \begin{cases} (1-\gamma)(\mathbf{e}_1^{L+2})^T \boldsymbol{\epsilon}_{t-1} & i = 0, \\ \gamma(\mathbf{e}_1^{L+2})^T \boldsymbol{\epsilon}_{t-1} + (\mathbf{e}_{L+3})^T \boldsymbol{\epsilon}_{t-1} & i = 1, \\ (\mathbf{e}_{L+i+2})^T \boldsymbol{\epsilon}_{t-1} & i \geq 2. \end{cases}$$

From the above expression, we can clearly see that the effect of the optimal weight vector is essentially to postpone a fraction γ of the order quantity $(\mathbf{e}_1^{L+2})^T \boldsymbol{\epsilon}_{t-1}$ from the current period to the subsequent period, where $(\mathbf{e}_1^{L+2})^T \boldsymbol{\epsilon}_{t-1}$ is the optimal order quantity when the retailer optimizes its own cost (see Proposition 1). This policy is very easy to implement. We note that the reason the postponement is only made to the subsequent period is because the supplier's lead time is one period in this case.

From the formula of γ, we observe that when the total supply chain costs are optimized, γ is always greater than zero. By postponing a fraction γ of order revision quantity, the order uncertainty faced by the supplier is effectively reduced. When the retailer optimizes its own cost, the resulting order revision quantity is essentially the case of $o_{t,\,t+i}^*$ with $\gamma = 0$, which gives a maximum level of order uncertainty for the supplier. Therefore, the parameter γ can be viewed as a control level that regulates the supplier's order uncertainty: when γ increases from 0 to 1, the supplier's order uncertainty decreases from the maximum level to zero.

From Proposition 4, we observe that if $(h + p) = (h' + p')$, or $k = 1$, then we have $\gamma = 1$ for sure. This means that when the retailer and the supplier have the same cost parameters, to minimize the total supply chain costs, the retailer should postpone the entire current-period order revision quantity to the subsequent period. As a result, the supplier faces no order uncertainty and no safety stock is needed. The total supply chain costs thus only constitute the retailer's cost.

Proposition 5. *The optimal γ is decreasing in the relative cost ratio k and increasing in the relative uncertainty ratio v. The resulting order uncertainty for the supplier $\Delta_o(\mathbf{w}_1^*)$ is decreasing in γ.*

The insights obtained from the above proposition are worth further comment. If the retailer's cost parameters are relatively higher than the supplier's cost parameters, then a smaller fraction of the order should be postponed. In this case, the supplier is left with more order uncertainty, which is intuitive. However, if the retailer's demand uncertainty becomes relatively higher than the supplier's order uncertainty, the proposition indicates that a larger fraction of the order should be postponed. This seems counter intuitive. The explanation is that when the retailer has higher demand uncertainty than the supplier's order uncertainty, it would cost less for the retailer to add additional safety stock than for the supplier to do so, because the incurred safety stock cost is concave increasing in the uncertainty level (more precisely, a square root relationship). Therefore, the retailer should postpone a larger fraction of the order to reduce the supplier's order uncertainty.

Numerical illustrations of the relationship between γ and k, v are given in Figures 2 and 3. Figure 2 shows the cases where the relative cost ratio k varies from 1 to 5 and the relative uncertainty ratio v is kept at fixed values of 1, 3, and 5. As we can see from the figure, when the relative cost ratio k increases, the optimal γ decreases. It can be also observed from the figure that when the relative uncertainty ratio v increases, the optimal γ increases. The same effect is further illustrated in Figure 3, where the value of v is varied from 1 to 15 and k is kept at fixed values of 2, 3, and 4.

5.1. The Bullwhip Effect

Before proceeding to discuss the bullwhip effect implications, we would like to comment that our setting of sharing order projections enables us to sharply

Figure 2 **Impact of Relative Cost Ratio on** γ ($v = 1, 3, 5$)

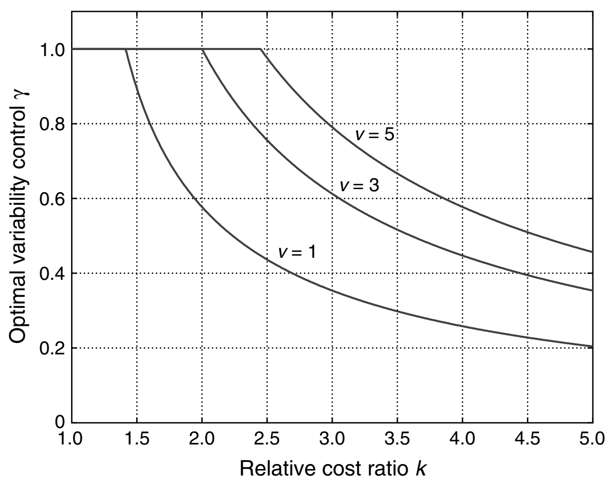

delineate between the effect of order variability and the effect of order uncertainty. It is fine that, in the traditional model setting without order projection, one focuses on order variability (the bullwhip effect metric), as that would be the key driver of cost. But when order projections are given, order variability is not the key driver any more. It is the uncertainty of the order revision that constitutes the cost driver to the supplier. This is why, even if the order variability increases, if there is less uncertainty in the order revision, the supplier can be better off. With this in mind, let us examine the bullwhip effect under the optimal weight $\mathbf{w}_i^*$ that minimizes the total supply chain costs. The total order variability in this case can be shown as, for $0 < \gamma \le 1$,

$$\mathrm{var}(O_t) = \sigma^2 + [(1-\gamma)^2 + \gamma^2](\mathbf{e}_1^{L+2})^T \boldsymbol{\Sigma} \mathbf{e}_1^{L+2} + 2\gamma(\mathbf{e}_1^{L+2})^T \boldsymbol{\Sigma} \mathbf{e}_{L+3} - \sum_{i=1}^{L+2} (\mathbf{e}_i)^T \boldsymbol{\Sigma} \mathbf{e}_i.$$

Figure 3 **Impact of Relative Uncertainty Ratio on** γ ($k = 2, 3, 4$)

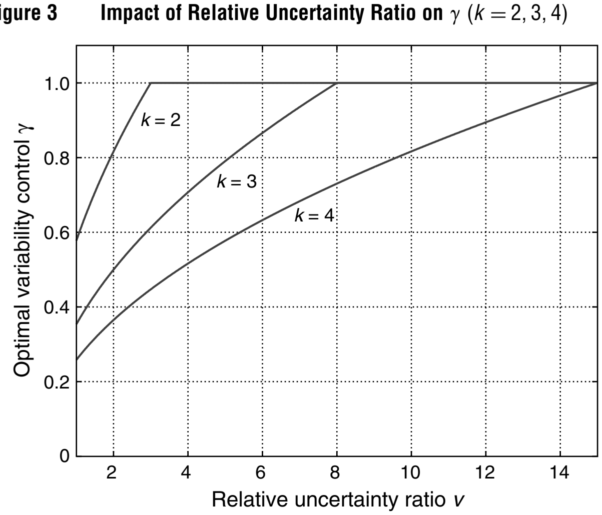

Therefore, the bullwhip effect result can be summarized as follows:

Proposition 6. *Under the optimal weight* $\mathbf{w}_i^*$, *the order variability amplification ratio is*

$$\frac{\mathrm{var}(O_t)}{\mathrm{var}(D_{t-1})} = 1 + \frac{1}{\sigma^2}\left([(1-\gamma)^2 + \gamma^2](\mathbf{e}_1^{L+2})^T \boldsymbol{\Sigma} \mathbf{e}_1^{L+2} + 2\gamma(\mathbf{e}_1^{L+2})^T \boldsymbol{\Sigma} \mathbf{e}_{L+3} - \sum_{i=1}^{L+2} (\mathbf{e}_i)^T \boldsymbol{\Sigma} \mathbf{e}_i \right), \quad (22)$$

with $0 < \gamma \le 1$. *Furthermore, the variability amplification ratio satisfies the following*:

(1) *If* $(\mathbf{e}_1^{L+2})^T \boldsymbol{\Sigma} \mathbf{e}_{L+3} \le -(\mathbf{e}_1^{L+2})^T \boldsymbol{\Sigma} \mathbf{e}_1^{L+2}$, *the variability amplification ratio is decreasing in* γ.

(2) *If* $(\mathbf{e}_1^{L+2})^T \boldsymbol{\Sigma} \mathbf{e}_{L+3} \ge (\mathbf{e}_1^{L+2})^T \boldsymbol{\Sigma} \mathbf{e}_1^{L+2}$, *the variability amplification ratio is increasing in* γ.

(3) *If* $|(\mathbf{e}_1^{L+2})^T \boldsymbol{\Sigma} \mathbf{e}_{L+3}| < (\mathbf{e}_1^{L+2})^T \boldsymbol{\Sigma} \mathbf{e}_1^{L+2}$, *the variability amplification ratio is convex in* γ, *with the minimum attained at* $\gamma = 1/2 - (\mathbf{e}_1^{L+2})^T \boldsymbol{\Sigma} \mathbf{e}_{L+3} / 2(\mathbf{e}_1^{L+2})^T \boldsymbol{\Sigma} \mathbf{e}_1^{L+2}$.

We know that when the retailer optimizes its own cost, it is equivalent to $\gamma = 0$. So the order variability amplification ratio (22) can be viewed as a further generalization of the bullwhip effect result obtained in Proposition 1. Let us consider a few examples.

First, if $\gamma = 1$ (e.g., in the case of $k = 1$), the supplier would face no order uncertainty because of the complete order postponement by the retailer. But the total order variability observed by the supplier might still be high. The order variability amplification effect in this case is determined by the sign of $(\mathbf{e}_1^{L+2})^T \boldsymbol{\Sigma} \mathbf{e}_1^{L+2} + 2(\mathbf{e}_1^{L+2})^T \boldsymbol{\Sigma} \mathbf{e}_{L+3} - \sum_{i=1}^{L+2} (\mathbf{e}_i)^T \boldsymbol{\Sigma} \mathbf{e}_i$. When the middle term is positive, the order variability amplification effect is in fact greater than if $\gamma = 0$ (i.e., with no order postponement). So from this example, we see that increasing γ only serves to reduce the order uncertainty for the supplier, but not necessarily the total order variability.

Another interesting example is the case when forecast information is not available beyond the lead time L, which is fairly common in practice. Effectively we have $(\mathbf{e}_1^{L+2})^T \boldsymbol{\Sigma} \mathbf{e}_{L+3} = 0$. By Proposition 6 part (3), the minimum order variability amplification ratio is attained at $\gamma = 1/2$. When γ increases from 1/2 to 1, the order variability is in fact increasing.

To give a visual illustration of the relationship between the order variability amplification ratio and γ for the three cases described in Proposition 6, we choose the following values for the relevant parameters: $\sigma^2 = 1$, $(\mathbf{e}_1^{L+2})^T \boldsymbol{\Sigma} \mathbf{e}_1^{L+2} = 0.2$, $\sum_{i=1}^{L+2} (\mathbf{e}_i)^T \cdot \boldsymbol{\Sigma} \mathbf{e}_i = 0.1$, and $(\mathbf{e}_1^{L+2})^T \boldsymbol{\Sigma} \mathbf{e}_{L+3} = \pm 0.5, 0$, where the values chosen for the terms $(\mathbf{e}_1^{L+2})^T \boldsymbol{\Sigma} \mathbf{e}_1^{L+2}$ and $\sum_{i=1}^{L+2} (\mathbf{e}_i)^T \boldsymbol{\Sigma} \mathbf{e}_i$ ensure positive correlation across forecast revision quantities over the exposure period, and the three values for $(\mathbf{e}_1^{L+2})^T \boldsymbol{\Sigma} \mathbf{e}_{L+3}$ are chosen to illustrate the three cases given in Proposition 6. The results are plotted in Figure 4.

Figure 4 **Order Variability Amplification Ratio vs.** γ ($x = (\mathbf{e}_1^{L+2})^T \mathbf{\Sigma} \mathbf{e}_{L+3} = \pm 0.5, 0$)

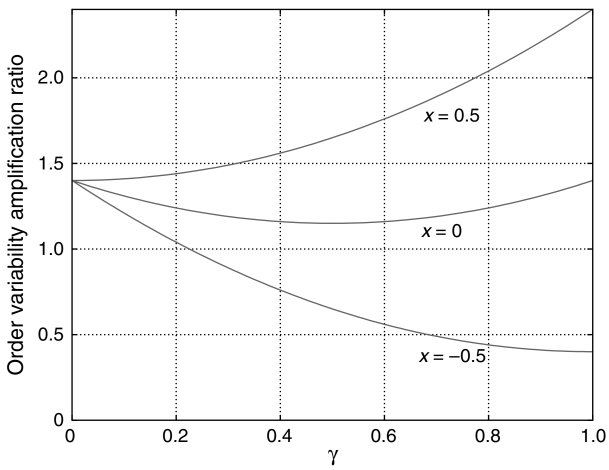

Figure 4 illustrates the fact that when γ increases (i.e., the order uncertainty faced by the supplier decreases), the manifested bullwhip effect in terms of total order variability in some cases would actually increase. This is because the total order variability is determined not only by the reduced order uncertainty (due to postponement), but also by the cross correlation between the postponed quantity and the subsequent-period order quantity; and the latter is not necessarily decreasing, depending on the characteristics of the underlying MMFE demand process.

This observation calls for a further discussion on the difference between the effect of uncertainty and the effect of variability in supply chains. When demand is i.i.d. with no advance demand signals, the concepts of uncertainty and variability are equivalent (see §2). However, when demand is more complex, the demand uncertainty is the conditional variance given the available information, which is different from the unconditional demand variance (i.e., the total variability of demand). This distinction was also discussed by Aviv (2001) in an i.i.d. demand model with MMFE market signals. He showed that information sharing can help to reduce the order uncertainty, but the order variability would remain the same. In our model, we have shown that information sharing coupled with order postponement can actually increase the order variability, even though the order uncertainty is reduced.

In general, the demand/order uncertainty affects the inventory-related supply chain cost as safety stock is required to buffer the risk. Yet the demand/order variability is usually associated with the capacity-related supply chain cost. For example, additional shipping capacity cost may be incurred when the order quantity exceeds a certain full-truckload threshold (see Balakrishnan et al. 2004). In our cost model, we assumed that the shipping capacity is unlimited and thus that only the inventory-related cost is relevant. Therefore, it is not surprising that the resulting optimal policy reduces the order uncertainty, but not necessarily the total order variability. By introducing additional capacity-related cost into our model framework, we can further extend our results to balance the trade-off between inventory and capacity considerations. This will be discussed as a potential extension in the next section.

Moreover, the insights we obtained in this section complement those found by Aviv (2007). In a model setting symmetrical to ours, Aviv (2007) considered the case when the supplier is the sole observer of the demand signals. His findings suggest that when the supplier shares the market signals with the retailer, the supplier faces *higher* order uncertainty but *less* order variability, which is symmetrical to our results. He further suggested that it would be better for the supply chain if the supplier not only maintained a low service level, but practically eliminated all safety stock inventories. It is interesting to point out that his recommendation achieves the same effect as in our case with $\gamma = 1$, where the order uncertainty is completely eliminated and thus no safety stock is needed at the supplier side.

In summary, we have derived under a very general setting the optimal order-smoothing weight to minimize the total supply chain costs. The result has a surprisingly simple and elegant form that bears similarity to the postponement strategy. We have shown the relationship between the optimal control and its two main drivers: the relative cost ratio and the relative uncertainty ratio between the retailer and the supplier. Based on the optimal control, we derive the general formula for the order variability amplification ratio. We then discuss the difference between the order uncertainty and order variability and explain why reducing the order uncertainty in our case does not necessarily ensure the dampening of the order variability (the bullwhip effect). In other words, a coordinated supply chain does not have to exhibit minimum bullwhip effect; instead, the manifested bullwhip effect will vary, depending on the cost structure and the relative uncertainty/information structure of the supply chain.

6. Extensions and Concluding Remarks

The results derived in §§4 and 5 are based on the assumption that the supplier's production lead time $l = 1$. As discussed in §3.3, when $l > 1$, exact analysis of the supplier's exposure-period demand distribution becomes cumbersome, as it resembles a lost-sales system. In this case, a simplification approach is to assume that the shortfall inventory is "borrowed" from the secondary source, which needs to be returned to the secondary source after usage (see Lee et al. 2000). This assumption effectively turns the supplier's inventory problem into a full-backlog problem. Therefore, analysis of the supplier's exposure-period demand distribution becomes straightforward. We present results with $l > 1$ under this simplifying assumption.

Recall that the retailer's order quantity is given by (4). With $l > 1$, under the "borrowed" inventory assumption the supplier's long-run average cost is given by

$$\begin{aligned} C'(m', \mathbf{w}'_1, \mathbf{w}'_2, \ldots, \mathbf{w}_1, \mathbf{w}_2, \ldots) &= E\left\{c'\left(S'_t - \sum_{i=1}^{l} O_{t+i}\right)\right\} \\ &= E\left\{c'\left(m' + \sum_{i=1}^{\infty} \mathbf{w}'^T_i \boldsymbol{\epsilon}_{t-i} - l\mu - \sum_{i=1}^{l}\sum_{j=1}^{\infty}(\mathbf{w}_j - \mathbf{w}_{j-1} + \mathbf{e}_j)^T \boldsymbol{\epsilon}_{t+i-j}\right)\right\} \\ &= E\left\{c'\left(m' - l\mu + \sum_{i=1}^{\infty}(\mathbf{w}'_i - \mathbf{w}_{i+l} + \mathbf{w}_i - \mathbf{e}_{i+1}^{i+l})^T \boldsymbol{\epsilon}_{t-i} - \sum_{i=1}^{l}(\mathbf{w}_i + \mathbf{e}_1^i)^T \boldsymbol{\epsilon}_{t+l-i}\right)\right\}. \end{aligned}$$

Analogous to the definitions (9) and (10), let us define the following terms:

$$\begin{aligned} \Delta_f^l(\mathbf{w}_1, \mathbf{w}_2, \ldots) &= \mathrm{var}\left(\sum_{i=1}^{\infty}(\mathbf{w}_{i+l} - \mathbf{w}_i + \mathbf{e}_{i+1}^{i+l})^T \boldsymbol{\epsilon}_{t-i}\right) \\ &= \sum_{i=1}^{\infty}(\mathbf{w}_{i+l} - \mathbf{w}_i + \mathbf{e}_{i+1}^{i+l})^T \boldsymbol{\Sigma}(\mathbf{w}_{i+l} - \mathbf{w}_i + \mathbf{e}_{i+1}^{i+l}), \quad \text{and} \end{aligned}$$

$$\begin{aligned} \Delta_o^l(\mathbf{w}_1, \ldots, \mathbf{w}_l) &= \mathrm{var}\left(\sum_{i=1}^{l}(\mathbf{w}_i + \mathbf{e}_1^i)^T \boldsymbol{\epsilon}_{t+l-i}\right) \\ &= \sum_{i=1}^{l}(\mathbf{w}_i + \mathbf{e}_1^i)^T \boldsymbol{\Sigma}(\mathbf{w}_i + \mathbf{e}_1^i). \end{aligned}$$

Proposition 7. *If the retailer shares with the supplier the forecast revision information* $\boldsymbol{\epsilon}_t$ *and the order weight vector* $\mathbf{w}_i$ $(i > 0)$, *then the net cost savings for the supplier are given by*

$$(h' + p')\phi(z')\left(\sqrt{\Delta_f^l(\mathbf{w}_1, \mathbf{w}_2, \ldots) + \Delta_o^l(\mathbf{w}_1, \ldots, \mathbf{w}_l)} - \sqrt{\Delta_o^l(\mathbf{w}_1, \ldots, \mathbf{w}_l)}\right).$$

This proposition extends Proposition 2. Analogously, we can obtain the following result for optimal order variability control when the supplier's production lead time is l periods:

Proposition 8. *The optimal weight vector* $\mathbf{w}_i^*$ *that minimizes the total supply chain costs is*

$$\mathbf{w}_i^* = \begin{cases} -\gamma \mathbf{e}_1^i + (1-\gamma)\mathbf{e}_{i+1}^{L+i+1} & 1 \le i \le l, \\ \mathbf{e}_{i+1}^{i+L+1} & i \ge l+1, \end{cases}$$

where γ *is given by*

$$\gamma = \begin{cases} \sqrt{\dfrac{v}{k^2 - 1}} & \text{if } k > \sqrt{1+v}, \\ 1 & \text{if } k \le \sqrt{1+v}, \end{cases}$$

with k *defined in* (19) *and* $v = \Delta_d / \Delta_o^l(\mathbf{e}_2^{L+2}, \ldots, \mathbf{e}_{l+1}^{L+l+1})$.

Substituting the optimal weight vector $\mathbf{w}_i^*$ into the retailer's future order revision quantity $o_{t,\,t+i}$ as specified in Proposition 3, we derive that

$$o_{t,\,t+i}^* = \begin{cases} (1-\gamma)(\mathbf{e}_1^{L+2})^T \boldsymbol{\epsilon}_{t-1} & i = 0, \\ (1-\gamma)(\mathbf{e}_{L+i+2})^T \boldsymbol{\epsilon}_{t-1} & 1 \le i \le l-1, \\ \gamma(\mathbf{e}_1^{L+l+1})^T \boldsymbol{\epsilon}_{t-1} + (\mathbf{e}_{L+l+2})^T \boldsymbol{\epsilon}_{t-1} & i = l, \\ (\mathbf{e}_{L+i+2})^T \boldsymbol{\epsilon}_{t-1} & i \ge l+1. \end{cases}$$

From the expression, we can see that when the supplier's production lead time is l periods, the effect of the optimal weight vector is essentially to postpone a fraction γ of all order revision quantities (determined when the retailer optimizes its own cost) from the current period and the next $l-1$ periods to the period l. Therefore, Proposition 8 generalizes the postponement insight we have obtained in Proposition 4.

We showed in §4 that the order-revision process is also an MMFE process. Under the "borrowed" inventory assumption, generalizing this result from the two-stage setting to a more complex supply chain network is straightforward. With each supply chain node facing an MMFE demand process, our analysis of the two-stage supply chain can then be extended to the entire supply chain network.

In this paper, we have considered the information sharing and order variability control problem, assuming that the retailer is willing to minimize the total supply chain costs. Here we discuss a few potential extensions to address the incentive issue. First, as we discussed in §5.1, we can extend our model to consider order variability as a source of retailer's cost. For example, additional shipping capacity cost may be incurred to the retailer when its order exceeds a certain full-truckload threshold. In addition, we can also assume that the retailer incurs a unit cost for receiving orders from the secondary source. This cost represents the inconvenience and disrupted workflow (e.g., irregular receiving hours) of handling orders arriving from the secondary source. We can further assume that a unit transfer cost is incurred by the supplier to reimburse the retailer's inconvenience cost. By designing the transfer cost contract between the retailer and the supplier, a desired level of coordination can be achieved between the two parties. In this case, the analysis as well as the insight would be similar to the derivation of Propositions 4–8. Second, we can address the incentive issue by directly designing a contract mechanism to achieve the first-best integrated system performance, such as Lee and Whang (1999) and Cachon and Lariviere (2001). This direction merits further research.

A key observation made in this research is about the implementation issue of information sharing between the retailer and the supplier. We argue that sharing the retailer's projected future orders is a more effective way than just sharing the POS and/or forecast information. By generating projected order information for the supplier, the retailer not only fulfills the task of sharing the demand information but also eliminates the need for the supplier to guess the retailer's underlying demand model and order policy. All the supplier needs to do is to focus on the historical projected order information to determine the stochastic characterization of the order revision process, which is much easier than the task of reproducing the retailer's forecasting and ordering logic.

Another key observation is that, by sharing order projections, we can have a setting that enables us to sharply delineate between the effect of order variability and the effect of order uncertainty. In the traditional model setting without order projection, order variability (same as order uncertainty) is the key cost driver. But in the case when order projections are given, order variability is not the key driver any more. It is the uncertainty of the order revision that constitutes the cost driver to the supplier. In this sense, we are also bringing in an intriguing future research question. So far, most researchers, including Cachon et al. (2007), have been looking at order variability as the measure of the bullwhip effect. Maybe we need to develop a new measure of the harmful effect of the bullwhip effect, i.e., a measure that captures the order uncertainty and not just the variability.

Finally, in an experimental study, Croson and Donohue (2003) showed that sharing POS data helps reduce the bullwhip effect in the Beer Game. It would be interesting to see what happens if order projections are shared within the Beer Game. With each node of the Beer Game supply chain sharing order projections, we believe the bullwhip effect can be further reduced. The information delay problem in the Beer Game can also be easily handled in our model analysis—just shift the order projections with the appropriate information lead time. Chen (1999) showed that the effect of information lead time is the same as that of transportation/production lead time in determining the optimal decision rules. This effect holds true in our model framework too, except that the value of information sharing may diminish as information lead time increases. But the supply chain would still be better off than the case when no information is shared.

Acknowledgments
The authors thank the associate editor and two anonymous referees for their valuable comments and suggestions. They also thank the seminar participants at Washington University, The Wharton School, Dartmouth College, and Duke University for helpful discussions.

Appendix

Proof of Proposition 1. By definition (2), the retailer's exposure-period demand can be expressed as $\sum_{i=0}^{L} D_{t+i} = \sum_{i=0}^{L} F_{t-1,\,t+i} + \sum_{i=0}^{L}\sum_{j=0}^{i} \epsilon_{t+i-j,\,t+i}$. Given the free-return assumption, the retailer's problem is a standard newsvendor problem. So the optimal order-up-to level is given by $S_t = \sum_{i=0}^{L} F_{t-1,\,t+i} + z\sqrt{\Delta_d} = \sum_{i=1}^{\infty} (\mathbf{e}_{i+1}^{i+L+1})^T \boldsymbol{\epsilon}_{t-i} + (L+1)\mu + z\sqrt{\Delta_d}$, where $z = \Phi^{-1}(p/(h+p))$, Δ_d is defined in (5), and the second equality follows from the definition of (2). So S_t in this case is an affine time-invariant function of forecast revision vectors $\{\boldsymbol{\epsilon}_{-\infty}, \ldots, \boldsymbol{\epsilon}_{t-1}\}$ with the weight vector $\mathbf{w}_i^* = \mathbf{e}_{i+1}^{i+L+1}$ for $i > 0$ and the constant $m^* = (L+1)\mu + z\sqrt{\Delta_d}$. The order variability amplification ratio follows from substituting $\mathbf{w}_i^*$ into expression (4) and then assessing the variance of O_t. □

Proof of Propositions 2 and 3. Follow the derivation preceding the proposition. □

Proof of Proposition 4. To minimize (18), we can first set $\mathbf{w}_i = \mathbf{e}_{i+1}^{i+L+1}$ for $i \geq 2$ according to the definition of $\Delta_w(\mathbf{w}_1, \mathbf{w}_2, \ldots)$. So the problem is reduced to search for the optimal $\mathbf{w}_1$ to minimize

$$f(\mathbf{w}_1) = (h+p)\phi(z)\sqrt{(\mathbf{w}_1 - \mathbf{e}_2^{L+2})^T \boldsymbol{\Sigma} (\mathbf{w}_1 - \mathbf{e}_2^{L+2}) + \Delta_d} + (h'+p')\phi(z')\sqrt{(\mathbf{w}_1 + \mathbf{e}_1)^T \boldsymbol{\Sigma} (\mathbf{w}_1 + \mathbf{e}_1)}.$$

From the above expression, it is easy to see that an optimal vector $\mathbf{w}_1$ will have zeros for all elements with an index greater than $L+2$. Therefore, we only need to consider the problem with a truncated $\mathbf{w}_1$. Below we will use the same notation for derivation with the understanding that all vectors are truncated to the first $(L+2)$ elements and $\boldsymbol{\Sigma}$ is the first $(L+2) \times (L+2)$ truncation of the original matrix. Take the gradient of $f(\mathbf{w}_1)$ with respect to the truncated $\mathbf{w}_1$. We have

$$\nabla f = (h+p)\phi(z)\frac{(\mathbf{w}_1 - \mathbf{e}_2^{L+2})^T \boldsymbol{\Sigma}}{\sqrt{(\mathbf{w}_1 - \mathbf{e}_2^{L+2})^T \boldsymbol{\Sigma} (\mathbf{w}_1 - \mathbf{e}_2^{L+2}) + \Delta_d}} + (h'+p')\phi(z')\frac{(\mathbf{w}_1 + \mathbf{e}_1)^T \boldsymbol{\Sigma}}{\sqrt{(\mathbf{w}_1 + \mathbf{e}_1)^T \boldsymbol{\Sigma} (\mathbf{w}_1 + \mathbf{e}_1)}}.$$

Note that ∇f does not exist at $\mathbf{w}_1 = -\mathbf{e}_1$. Let us first examine $f(\mathbf{w}_1)$ within the neighborhood of $\mathbf{w}_1 = -\mathbf{e}_1 + \delta\mathbf{u}$, where δ is a scalar and $\mathbf{u}$ is an arbitrary unitary vector. It is easy to show that

$$f(-\mathbf{e}_1+\delta\mathbf{u})-f(-\mathbf{e}_1)\geq (h+p)\phi(z)\frac{-\mathbf{u}^T\boldsymbol{\Sigma}\mathbf{e}_1^{L+2}}{\sqrt{(\mathbf{e}_1^{L+2})^T\boldsymbol{\Sigma}\mathbf{e}_1^{L+2}+\Delta_d}}\cdot\delta + (h'+p')\phi(z')\sqrt{\mathbf{u}^T\boldsymbol{\Sigma}\mathbf{u}}\cdot|\delta|.$$

Because $\boldsymbol{\Sigma}$ is semipositive definite, by Cauchy-Schwartz inequality, we have

$$|\mathbf{u}^T\boldsymbol{\Sigma}\mathbf{e}_1^{L+2}|\leq\sqrt{(\mathbf{e}_1^{L+2})^T\boldsymbol{\Sigma}\mathbf{e}_1^{L+2}}\cdot\sqrt{\mathbf{u}^T\boldsymbol{\Sigma}\mathbf{u}}.$$

Let $k=(h+p)\phi(z)/(h'+p')\phi(z')$; we have

$$f(-\mathbf{e}_1+\delta\mathbf{u})-f(-\mathbf{e}_1) \geq (h'+p')\phi(z')\left[k\frac{\operatorname{sgn}\{\mathbf{u}^T\boldsymbol{\Sigma}\mathbf{e}_1^{L+2}\cdot\delta\}}{\sqrt{(\mathbf{e}_1^{L+2})^T\boldsymbol{\Sigma}\mathbf{e}_1^{L+2}+\Delta_d}}+\frac{1}{\sqrt{(\mathbf{e}_1^{L+2})^T\boldsymbol{\Sigma}\mathbf{e}_1^{L+2}}}\right] \cdot|\mathbf{u}^T\boldsymbol{\Sigma}\mathbf{e}_1^{L+2}\cdot\delta|.$$

Therefore, if $k\leq\sqrt{1+\Delta_d/(\mathbf{e}_1^{L+2})^T\boldsymbol{\Sigma}\mathbf{e}_1^{L+2}}$, we have $f(-\mathbf{e}_1+\delta\mathbf{u})-f(-\mathbf{e}_1)\geq 0$ for all δ and $\mathbf{u}$. Hence, in this case the optimal solution is $\mathbf{w}_1^*=-\mathbf{e}_1$.

Now for the case with $k>\sqrt{1+\Delta_d/(\mathbf{e}_1^{L+2})^T\boldsymbol{\Sigma}\mathbf{e}_1^{L+2}}$: by solving the first-order necessary condition $\nabla f=\mathbf{0}$, we obtain $\mathbf{w}_1^*=-\gamma\mathbf{e}_1+(1-\gamma)\mathbf{e}_2^{L+2}$, with

$$\gamma=\frac{\sqrt{(\mathbf{w}_1^*-\mathbf{e}_2^{L+2})^T\boldsymbol{\Sigma}(\mathbf{w}_1^*-\mathbf{e}_2^{L+2})+\Delta_d}}{k\sqrt{(\mathbf{w}_1^*+\mathbf{e}_1)^T\boldsymbol{\Sigma}(\mathbf{w}_1^*+\mathbf{e}_1)}+\sqrt{(\mathbf{w}_1^*-\mathbf{e}_2^{L+2})^T\boldsymbol{\Sigma}(\mathbf{w}_1^*-\mathbf{e}_2^{L+2})+\Delta_d}}.$$

Substituting $\mathbf{w}_1^*$ into the γ expression, we obtain $\gamma=\sqrt{\Delta_d/(k^2-1)(\mathbf{e}_1^{L+2})^T\boldsymbol{\Sigma}\mathbf{e}_1^{L+2}}=\sqrt{v/(k^2-1)}$. The optimality of $\mathbf{w}_1^*$ can be further verified by the Hessian matrix at $\mathbf{w}_1^*$. □

Proof of Proposition 5. Use the γ expression of Proposition 4 and the fact that $\Delta_o(\mathbf{w}_1^*)=(\mathbf{w}_1^*+\mathbf{e}_1)^T\boldsymbol{\Sigma}(\mathbf{w}_1^*+\mathbf{e}_1)=(1-\gamma)^2(\mathbf{e}_1^{L+2})^T\boldsymbol{\Sigma}\mathbf{e}_1^{L+2}$, which is decreasing in γ for $0<\gamma\leq 1$. □

Proof of Proposition 6. Because the total demand variability is σ^2, the variability amplification ratio (22) is obtained by dividing $\operatorname{var}(O_t)$ by σ^2. By taking the first- and second-order derivative of the expression (22) with respect to γ, we obtain the sensitivity results given in parts (1)–(3). □

Proof of Proposition 7. Follow the similar derivation steps of Proposition 2. □

Proof of Proposition 8. The problem is to find $\mathbf{w}_i$ to minimize the following function:

$$(h+p)\phi(z)\sqrt{\Delta_w(\mathbf{w}_1,\mathbf{w}_2,\ldots)+\Delta_d} + (h'+p')\phi(z')\sqrt{\Delta_o^l(\mathbf{w}_1,\ldots,\mathbf{w}_l)}.$$

To do this, we can first set $\mathbf{w}_i=\mathbf{e}_{i+1}^{L+i+1}$ for $i\geq l+1$ according to the definition of $\Delta_w(\mathbf{w}_1,\mathbf{w}_2,\ldots)$. Therefore, the problem is reduced to search for the optimal $\mathbf{w}_1,\ldots,\mathbf{w}_l$ that minimizes

$$f(\mathbf{w}_1,\ldots,\mathbf{w}_l) =(h+p)\phi(z)\sqrt{\sum_{i=1}^{l}(\mathbf{w}_i-\mathbf{e}_{i+1}^{L+i+1})^T\boldsymbol{\Sigma}(\mathbf{w}_i-\mathbf{e}_{i+1}^{L+i+1})+\Delta_d} +(h'+p')\phi(z')\sqrt{\sum_{i=1}^{l}(\mathbf{w}_i+\mathbf{e}_1^i)^T\boldsymbol{\Sigma}(\mathbf{w}_i+\mathbf{e}_1^i)}.$$

From the above expression, it is easy to see that an optimal vector $\mathbf{w}_i$ $(1\leq i\leq l)$ will have zeros for all elements with indexes greater than $L+l+1$. Therefore, we only need to consider the problem with all vectors being truncated to the first $L+l+1$ elements and $\boldsymbol{\Sigma}$ being truncated to the first $(L+l+1)\times(L+l+1)$ elements. Below we will use the same notation with the understanding that all vectors are truncated to the first $(L+l+1)$ elements, and $\boldsymbol{\Sigma}$ is the first $(L+l+1)\times(L+l+1)$ truncation of the original matrix.

Now define the following vectors: $\mathbf{x}=[\mathbf{w}_1^T,\ldots,\mathbf{w}_l^T]^T$, $\mathbf{a}=[(\mathbf{e}_2^{L+2})^T,\ldots,(\mathbf{e}_{l+1}^{L+l+1})^T]^T$, and $\mathbf{b}=[(\mathbf{e}_1^1)^T,\ldots,(\mathbf{e}_1^l)^T]^T$. Also define the new $l(L+l+1)\times l(L+l+1)$ matrix $\mathbf{A}$ by stacking the $(L+l+1)\times(L+l+1)$ matrix $\boldsymbol{\Sigma}$ on the diagonal l times. Then the objective function can be rewritten as $f(\mathbf{x})=(h+p)\phi(z)\sqrt{(\mathbf{x}-\mathbf{a})^T\mathbf{A}(\mathbf{x}-\mathbf{a})+\Delta_d}+(h'+p')\cdot\phi(z')\sqrt{(\mathbf{x}+\mathbf{b})^T\mathbf{A}(\mathbf{x}+\mathbf{b})}$. This is the same problem as analyzed in Proposition 4. So by the result of Proposition 4, we can obtain the optimal solution $\mathbf{x}^*$. The result then follows from converting $\mathbf{x}^*$ back to $\mathbf{w}_1,\ldots,\mathbf{w}_l$. □

References

Aviv, Y. 2001. The effect of collaborative forecasting on supply chain performance. *Management Sci.* **47**(10) 1326–1343.

Aviv, Y. 2002. Gaining benefits from joint forecasting and replenishment processes: The case of auto-correlated demand. *Manufacturing Service Oper. Management* **4**(1) 55–74.

Aviv, Y. 2003. A time-series framework for supply chain inventory management. *Oper. Res.* **51**(2) 210–227.

Aviv, Y. 2007. On the benefits of collaborative forecasting partnerships between retailers and manufacturers. *Management Sci.* **53**(5) 777–794.

Balakrishnan, A., J. Geunes, M. S. Pangburn. 2004. Coordinating supply chains by controlling upstream variability propagation. *Manufacturing Service Oper. Management* **6**(2) 163–183.

Bourland, K. E., S. G. Powell, D. F. Pyke. 1996. Exploiting timely demand information to reduce inventories. *Eur. J. Oper. Res.* **92** 239–253.

Box, G. E. P., G. M. Jenkins, G. C. Reinsel. 1994. *Time Series Analysis: Forecasting and Control*, 3rd ed. Prentice-Hall, Englewood Cliffs, NJ.

Cachon, G., M. Fisher. 2000. Supply chain inventory management and the value of shared information. *Management Sci.* **46**(8) 1032–1048.

Cachon, G., M. Lariviere. 2001. Contracting to assure supply: How to share demand forecasts in a supply chain. *Management Sci.* **47**(5) 629–646.

Cachon, G., T. Randall, G. Schmidt. 2007. In search of the bullwhip effect. *Manufacturing Service Oper. Management* **9**(4) 457–479.

Chen, F. 1999. Decentralized supply chains subject to information delays. *Management Sci.* **45**(8) 1076–1090.

Chen, F. 2003. Information sharing and supply chain coordination. A. G. de Kok, S. C. Graves, eds. *Handbooks in Operations Research and Management Science, Vol. 11. Supply Chain Management*: *Design, Coordination, and Operation*. Elsevier, Amsterdam, 341–421.

Chen, F., Z. Drezner, J. K. Ryan, D. Simchi-Levi. 2000. Quantifying the bullwhip effect in a simple supply chain: The impact of forecasting, lead times, and information. *Management Sci.* **46**(3) 436–443.

Clark, T. H., J. Hammond. 1997. Reengineering channel reordering processes to improve total supply-chain performance. *Production Oper. Management* **6**(1) 248–265.

Croson, R., K. Donohue. 2003. Impact of POS data sharing on supply chain management: An experimental study. *Production Oper. Management* **12**(1) 1–11.

Gallego, G., O. Ozer. 2001. Integrating replenishment decisions with advance demand information. *Management Sci.* **47**(10) 1344–1360.

Gaur, V., A. Giloni, S. Seshadri. 2005. Information sharing in a supply chain under ARMA demand. *Management Sci.* **51**(6) 961–969.

Gavirneni, S., R. Kapuscinski, S. Tayur. 1999. Value of information in capacitated supply chains. *Management Sci.* **45**(1) 16–24.

Gilbert, K. 2005. An ARIMA supply chain model. *Management Sci.* **51**(2) 305–310.

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

Graves, S. C., D. B. Kletter, W. B. Hetzel. 1998. A dynamic model for requirements planning with application to supply chain optimization. *Oper. Res.* **46**(3) S35–S49.

Graves, S. C., H. C. Meal, S. Dasu, Y. Qiu. 1986. Two-stage production planning in a dynamic environment. S. Axsater, C. Schneeweiss, E. Silver, eds. *Multi-Stage Production Planning and Control. Lecture Notes in Economics and Mathematical Systems*, Vol. 266. Springer-Verlag, Berlin, 9–43.

Gullu, R. 1996. On the value of information in dynamic production/inventory problems under forecast evolution. *Naval Res. Logist.* **43** 289–303.

Gullu, R. 1997. A two-echelon allocation model and the value of information under correlated forecasts and demands. *Eur. J. Oper. Res.* **99** 386–400.

Hausman, W. H. 1969. Sequential decision problems: A model to exploit existing forecasters. *Management Sci.* **16**(2) B93–B111.

Heath, D. C., P. L. Jackson. 1994. Modelling the evolution of demand forecasts with application to safety stock analysis in production/distribution systems. *IIE Trans.* **26**(3) 17–30.

Iida, T., P. H. Zipkin. 2006. Approximate solutions of a dynamic forecast-inventory model. *Manufacturing Service Oper. Management* **8**(4) 407–425.

Johnson, G. D., H. E. Thompson. 1975. Optimality of myopic inventory policies for certain dependent demand processes. *Management Sci.* **21**(11) 1303–1307.

Kahn, J. A. 1987. Inventories and the volatility of production. *Amer. Econom. Rev.* **77** 667–679.

Kurt Salmon Associates. 1993. *Efficient Consumer Response*: *Enhancing Consumer Value in the Grocery Industry*. Food Marketing Institute, Washington, D.C.

Lee, H. L., S. Whang. 1999. Decentralized multi-echelon supply chains: Incentives and information. *Management Sci.* **45**(5) 633–640.

Lee, H. L., V. Padmanabhan, S. Whang. 1997. Information distortion in a supply chain: The bullwhip effect. *Management Sci.* **43**(4) 546–558.

Lee, H. L., K. C. So, C. S. Tang. 2000. The value of information sharing in a two-level supply chain. *Management Sci.* **46**(5) 626–643.

Liu, H., K.-L. Trui, F. Tsung. 2005. A nearly optimal order policy to reduce bullwhip effect. *Quality Tech. Quant. Management* **2**(1) 91–108.

Lu, X., J.-S. Song, A. Regan. 2006. Inventory planning with forecast updates: Approximate solutions and cost error bounds. *Oper. Res.* **54**(6) 1079–1097.

Miyaoka, J., W. H. Hausman. 2004. How a base stock policy using stale forecasts provides supply chain benefits. *Manufacturing Service Oper. Management* **6**(2) 149–162.

Porteus, E. 2002. *Foundations of Stochastic Inventory Theory*. Stanford University Press, Stanford, CA.

Raghunathan, S. 2001. Information sharing in a supply chain: A note on its value when demand is nonstationary. *Management Sci.* **47**(4) 605–610.

Terwiesch, C., Z. J. Ren, T. H. Ho, M. A. Cohen. 2005. An empirical analysis of forecast sharing in the semiconductor equipment supply chain. *Management Sci.* **51**(2) 208–220.

Toktay, L. B., L. M. Wein. 2001. Analysis of a forecasting-production-inventory system with stationary demand. *Management Sci.* **47**(9) 1268–1281.

Voluntary Interindustry Commerce Solutions Association. 1999. Roadmap to CPFR: The case studies. http://www.cpfr.org.