

Inventory Policies and Information Sharing: An Efficient Frontier Approach

René Caldentey

Booth School of Business, The University of Chicago, Chicago, IL 60637, rene.caldentey@chicagobooth.edu

Avi Giloni

Sy Syms School of Business, Yeshiva University, New York, NY 10033, agiloni@yu.edu

Clifford Hurvich

Stern School of Business, New York University, New York, NY 10012, churvich@stern.nyu.edu

We consider a two-tier inventory management system with one retailer and one supplier. The retailer serves a demand driven by a stationary moving average process (of possibly infinite order) and places periodic inventory replenishment orders to the supplier. In this setting, we study the value of information sharing and its impact on the retailer's optimal ordering strategy. We argue that information sharing affects performance through two key cost drivers: (i) on-hand inventory variability and (ii) replenishment order variability. We characterize a "*Pareto frontier*" between these two sources of variability by identifying optimal inventory replenishment strategies that trade-off one type of variability for the other in a cost efficient way. For the case in which the retailer is able to share her complete demand history, we provide a full characterization of the efficient frontier, as well as of an optimal replenishment policy. On the other hand, when the retailer is not able (or willing) to share any demand information we provide a partial characterization of an optimal solution and show that information sharing does not always add value. We also show that the question of identifying conditions under which information sharing does offer value reduces to a delicate analysis of the invertibility (in a time series sense) of a specific stationary process.

Key words: inventory control; supply chain management; information sharing; order smoothing

1. Introduction

Overview: The implementation of efficient inventory replenishment policies, in the context of decentralized supply chain systems, relies heavily on the ability of the different supply chain members (e.g., retailers, distributors, manufacturers and suppliers) to forecast future demand and adjust production and inventory levels accordingly. Lack of demand information and poor visibility of inventory positions throughout the system lead to suboptimal performance in terms of operating costs (excessive inventory) and service level (in the form of unsatisfied demand and backorders).

There exists an extensive body of work in the supply chain management literature dedicated to measure the value of information sharing and to study the design of contractual agreements

to support the exchange of demand and inventory data among different supply chain members; see, for example, [Gavirneni et al. \(1999\)](#), [Lee et al. \(2000\)](#), [Aviv \(2001, 2007\)](#), [Chen \(2003\)](#), [Gaur et al. \(2005\)](#), [Giloni et al. \(2014\)](#) and references therein. The goal of this paper is to contribute to this stream of literature by shedding some light on how information sharing impacts performance through two key cost drivers: (i) on-hand inventory variability and (ii) replenishment order variability. In particular, our work aims at providing a simple characterization of the tension between these two forms of variability as well as the underlying mechanism that connects them.

At a high level, the approach that we take to analyze the value of information sharing in supply chains resembles the one used by [Graves et al. \(1998\)](#) in the context of material requirement planning (MRP) systems and is motivated by the fact that inventory replenishment policies have a *direct* and an *indirect* effect on the operating costs of a firm.

Direct Cost Effect: On one hand, a replenishment policy impacts the evolution of the physical flows of inventory stocks and therefore has a direct effect on a firm's holding and backordering costs. The problem of minimizing these direct inventory costs has been at the forefront of most of the traditional literature on inventory theory (see, for example, [Zipkin 2000](#)). As we will argue later, under some mild stationarity assumptions on demand and costs, we can use the volatility of a firm's on-hand inventory to quantify this direct cost effect.

Indirect Cost Effect: On the other hand, an inventory replenishment policy also impacts the cost structure of upstream suppliers that replenish these orders. Indeed, the empirical literature on supply chain management and the bullwhip effect (see [Lee et al. 1997, 2004](#) and references therein) argues that orders' volatility is one of the fundamental drivers of system inefficiencies and inventory costs; in the form of unplanned purchase of supplies, inefficient utilization of production capacity and overtime, excess warehousing and transportation, among many other factors*. Thus, one should expect that an increase in order volatility should translate into higher procurement and inventory costs in the long run.

It is our impression that the problem of minimizing these indirect inventory costs has received much less attention in the inventory management literature, possibly because they materialize gradually over time and therefore can be less visible and much harder to quantify. However, forward-looking firms should make an effort to internalize the relationship between ordering decisions and their suppliers' performance and costs. In establishing this connection we argue that what matters is not the volatility of the order process per se but rather the volatility of the forecast error of these

* By some estimates these inefficiencies can increase inventory costs by as much as 12.5% to 25%, see [Kurt Salomon Associates \(1993\)](#).

orders as computed by upstream supply chain members. The point here is that the inefficiencies mentioned above are not exclusively driven by the fluctuations of a firm's orders but rather they are triggered by the inability of its suppliers to anticipate these orders and optimize production and distribution accordingly; *i.e.*, what matters is the forecast error experienced by upstream suppliers as they forecast their own demands.

Motivated by the previous discussion, in this paper we address the following research questions:

- First, how do the volatilities of inventories and orders' forecasts depend on the inventory replenishment strategy of a firm?

- Second, is it possible to characterize a "*Pareto frontier*" between these two sources of variability by identifying optimal inventory replenishment strategies that trade-off one type of variability for the other in a cost efficient way? Furthermore, what is the structure of these optimal ordering policies?

- Finally, is it possible to quantify the value of information sharing and its impact on the ordering process of a firm and its inventory related costs?

Summary of Methodology and Results: We investigate these research questions in the context of a two-tier supply chain with a single retailer and a single supplier as depicted in Figure 1.




Figure 1 A two-tier supply inventory system.

In our model, we assume that the retailer observes a demand driven by a stationary and invertible moving average process D_t of possibly infinite order. We also assume that the retailer places periodic inventory replenishment orders O_t to the supplier and that these orders also admit a representation given by a stationary (not necessarily invertible) moving average process. Finally, the supplier replenishes his inventory placing production requests P_t to an uncapacitated manufacturing source. (We will use female pronouns to refer to the retailer and male pronouns to refer to the supplier.)

The focal point of this paper is to study how the retailer should place her replenishment orders O_t so as to minimize her long-term average inventory and procurement costs. To do this, we argue that the connection between O_t and the retailer's operating costs can be captured using two specific variability measures: (A detailed mathematical description of the model will be provided in Section 2)

- **INVENTORY VARIABILITY:** $\sigma_I^2 := \lim_{t \rightarrow \infty} \text{Var}(I_t^R | \mathcal{M}^D(t))$. This is the stationary variance of the retailer's on-hand inventory I_t^R conditional on $\mathcal{M}^D(t)$, the history of observed demand.
- **MEAN SQUARED FORECAST ERROR (MSFE):** $\sigma_S^2 := \text{Var}[O_{t+1} - \mathbb{E}[O_{t+1} | \mathcal{M}_S(t)] | \mathcal{M}_S(t)]$, where $\mathcal{M}_S(t)$ is the supplier's available information in period t . Depending on the degree of collaboration between the retailer and the supplier, $\mathcal{M}_S(t)$ can include the history of orders as well as additional demand information that the retailer is willing to share. Two extreme cases of interest are: (a) *no information sharing* in which $\mathcal{M}_S(t)$ is generated exclusively by $\{O_t\}$ and *full information sharing* in which $\mathcal{M}_S(t)$ is generated by both $\{O_t\}$ and $\{D_t\}$.

The notion that σ_I has a direct impact on the retailer's operating costs is rather intuitive as holding and backordering costs are directly impacted by the volatility of her inventory process. On the other hand, the indirect relationship between σ_S and the retailer's costs can be less evident and will be discussed in detail in Section 6 using the supply chain system depicted in Figure 1. In the meantime, just to help the reader get some intuition, consider the retailer's per-period procurement costs $w^R O_t$, where w^R is the per unit wholesale price that the supplier charges the retailer. If the planning horizon is sufficiently large, it is reasonable to assume that more volatile orders translate into higher operating costs for the supplier, which he will pass to the retailer in the form of a higher wholesale price w^R . Thus, a forward-looking retailer should make replenishment decisions taking into account the fact that the wholesale price that she pays is a function of the variability of her own orders, that is, $w^R = w^R(\sigma_S)$.

For reasons that will become apparent in our analysis, we will distinguish and study the cases with and without information sharing separately. For the case in which the retailer is able and willing to share demand information with the supplier, we provide a complete characterization of σ_S and σ_I and their relationship. In particular, we show that there exists a *Pareto Frontier* of pairs (σ_S^*, σ_I^*) that defines the efficient trade-off between these two sources of variability. The point is that if the retailer wants to reduce the variability of her on-hand inventory she would have to increase the volatility of her orders to the supplier (as measured by the MSFE induced by her orders), and vice versa. Figure 2 depicts an example for the case in which market demand is given by a sequence of iid normally distributed random variables.

The retailer's choice of a particular pair (σ_S^*, σ_I^*) depends on her specific preferences over these two sources of volatilities. So, for example, if direct inventory costs are the primary cost component then the retailer would like to minimize σ_I by selecting a strategy close to point 'M' in Figure 2. On the other hand, if the indirect costs induced by the mean forecast error are more significant then a replenishment strategy close to point 'P' would be more appropriate. (Note that in point 'P' the

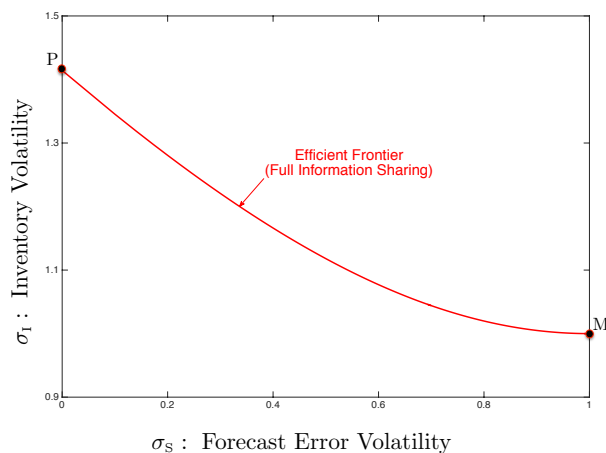


Figure 2 Locus of efficient volatility pairs (σ_S^*, σ_I^*) under full information sharing for the case in which market demand is driven by a sequence of iid normally distributed random variables with variance equal to one. (By Proposition 2 the Efficient frontier is independent of the mean demand.)

MSFE of the retailer's orders is exactly equal to zero which means that the supplier can perfectly predict the retailer's order in advance.) Of course, points 'M' and 'P' are only extreme cases and an optimal replenishment strategy needs to balance the tension between these two objectives.

We also analyze situations in which the retailer is unable (or unwilling) to share demand information with the supplier and show that information sharing does not always add value. We also show that the question of identifying conditions under which information sharing does offer value reduces to a delicate analysis of the invertibility (in a time series sense) of a specific stationary process. While the tools for checking invertibility numerically are available, a theoretical analysis of the value of information sharing has proven to be very challenging. At this point, we can only numerically observe and report a number of properties that an optimal replenishment policy should satisfy under no information sharing. Although more work needs to be done to provide a complete theoretical characterization of an optimal policy under no information sharing, we formulate a non-linear optimization problem that can be used to approximately compute a solution that has shown good numerical performance.

To illustrate some of the key takeaways of this paper, let us consider the example in Figure 3 that depicts the *efficient frontiers* associated with an optimal replenishment policy under full information sharing and no information sharing, and compare them to four order smoothing policies that have been proposed in the literature (a detailed description of these benchmark policies is provided in Section 5):

- BGP-MA: Under this policy the retailer's orders are equal to a moving average of the market demand (see Balakrishnan, Geunes, and Pangburn 2004).

- BGP-ES: In this case, the retailer's orders are computed using exponential smoothing of the market demand (see Balakrishnan, Geunes, and Pangburn 2004).
- GKH: Proposed by Graves, Kletter, and Hetzel (1998), this order smoothing policy is derived by minimizing the variance of the order process subject to a constraint that upper bounds the variance of the inventory process.
- CL: Finally, to quantify the effect of not internalizing the lack of invertibility of the retailer's orders, we consider a replenishment policy that minimizes cumulative supply chain (retailer and supplier) costs under the assumption that the supplier is always able to obtain the retailer's orders. We note that this policy coincides with the order smoothing policy proposed by Chen and Lee (2009) in those cases where the retailer's orders are indeed invertible.

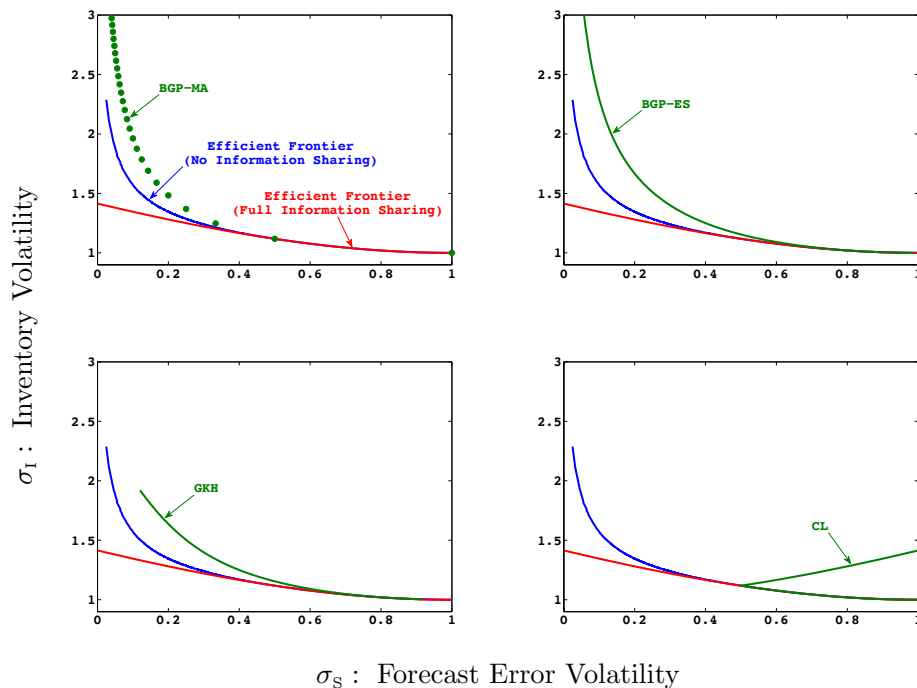


Figure 3 Each plot depicts the efficient frontiers between the forecast error (MSFE) and the inventory volatilities under full information sharing and no information sharing and compares them to one of the four benchmark policies: BGP-MA, BGP-ES, GKH and CL. In this example, market demand is iid and normally distributed.

In evaluating the performance of the four benchmark policies, we have assumed that the retailer does not share any demand information with the supplier. A first observation that emerges from the example in Figure 3 is that the efficient frontier under full information sharing dominates (i.e., produces simultaneously lower inventory and MSFE volatilities) the efficient frontier under no information sharing, in turn dominates the other four benchmark policies.

It is also worth noticing that the performance of the different policies appear to coincide in the region where $\sigma_s \geq 0.5$. In this region an efficient replenishment policy under full information sharing turns out to be an invertible process and so, as we will show later in Proposition 3, it is also implementable under no information sharing (see Example 2 in Section 4 for details). In other words, in this region, sharing demand information offers no real value. On the other hand, in the region where the retailer's efficient replenishment policy under full information sharing is not invertible, when $\sigma_s < 0.5$, information sharing can offer significant value. For example, under full information sharing the retailer can select an ordering policy that is fully predictable by the supplier (i.e., $\sigma_s = 0$) and at the same time keep the variability of her inventory process under control (i.e., $\sigma_I \approx 1.4142$). On the contrary, for all other policies, including an optimal strategy under no information sharing, if the retailer tries to select a predictable ordering policy then the variability of her inventory process will become unbounded.

The example in Figure 3 highlights the fundamental role that the notion of “invertibility” plays when (i) selecting an optimal inventory replenishment strategy and (ii) determining the value of information sharing. The goal of this paper is to shed some light on these issues by studying the mechanisms that connect replenishment decisions and inventory costs with demand forecasting information.

As a byproduct of our characterization of the efficient frontier between σ_s and σ_I , we are able to contrast two alternative measures of the bullwhip effect. On one hand, we use the traditional definition that computes the bullwhip as the ratio of order volatility to demand volatility, that is, $\text{Var}(O_{t+1})/\text{Var}(D_{t+1})$. We also consider an informationally-adjusted version in which the notion of variability that is used is the one that cannot be forecasted. In particular, we define this informationally-adjusted bullwhip measure by $\text{Var}(O_{t+1}|\mathcal{M}_S(t))/\text{Var}(D_{t+1}|\mathcal{M}_D(t))$. Using a set of numerical experiments, we show that these two measures can lead to significantly different recommendations on what an optimal replenishment policy should be. We argue, however, that the informationally-adjusted bullwhip provides a more accurate measure if we understand the bullwhip effect as a measure that connects the propagation of orders' variability and system inefficiencies.

We close this introduction by reviewing and positioning our results within the existing literature in supply chain management. We also include a brief subsection at the end with some definitions and notation.

1.1. Related Literature

In one of the first papers on information sharing in supply chains, Lee et al. (2000) consider a retailer that serves an AR(1) demand with a nonnegative autocorrelation coefficient, and places orders with

the manufacturer using a periodic-review order-up-to policy. In their model, both the manufacturer and the retailer know the parameters of the demand process; however, the retailer might choose not to share information about the actual realizations of demand with the manufacturer. Lee et al. (2000) analyze this model and conclude that information sharing results in significant inventory reduction and cost savings to the manufacturer. However, they qualify their results by noting that information sharing could be less valuable if the manufacturer uses the historical stream of orders from the retailer to forecast demand. Raghunathan (2001) further develops this idea and shows that the value of information sharing indeed decreases monotonically with each time period, converging to zero in the limit, under AR(1) demand with a nonnegative autocorrelation coefficient.

Zhang (2004) and Gaur et al. (2005) extend the original work of Lee et al. (2000) and Raghunathan (2001) by studying the value of information sharing in supply chains where the retailer serves an ARMA(p, q) demand as opposed to AR(1) demand. In each of these papers, the retailer places orders with a supplier using a periodic review order-up-to policy. Both the supplier and the retailer know the parameters of the demand process; however, the retailer may or may not choose to share information about the actual realizations of demand with the supplier. Zhang (2004) studies how the order process propagates upstream in a supply chain under the assumption that ARMA demand to the retailer and all upstream players is invertible. In such a case, there would be no value of information sharing to any of the players.

Zhang (2004) did not consider the case when demand becomes non-invertible (i.e., the current shock cannot be obtained as a linear combination of present and past demand observations only) at any stage of the chain, a phenomenon that Gaur et al. (2005) were first to point out can happen even though the retailer's demand is invertible. In other words, the supplier's demand may not be invertible with respect to the retailer's shocks, even though the retailer's demand is invertible with respect to its own shocks. Although the order process proposed by Gaur et al. (2005) can be used, their order process utilizes a suboptimal forecast as opposed to the best linear forecast of lead time demand and, hence, results in larger inventory-related costs compared with those under the best linear forecast as studied by Giloni et al. (2014).

Giloni et al. (2014) characterize where there is value to information sharing, describe those situations where there is potential benefit to the supplier if the retailer shares its demand information, determine when the maximum possible benefit may arise, and extend the characterization to further upstream players. They show that, in general, neither player $(k - 1)$'s order nor player k 's demand is necessarily an ARMA process with respect to their respective shocks, even though the retailer's demand is ARMA. Instead, demand propagates in general according to a quasi-ARMA, or QUARMA process, in which the most recent shock(s) may be absent.

In a series of papers, Aviv (2001, 2002, 2003, 2007) investigates demand propagation and information sharing in supply chains. In each of these papers, Aviv uses the Kalman filter to construct best linear forecasts and studies cases where the retailer faces demand that has a linear state-space representation. This includes stationary ARMA(p,q) demand models as well as more general non-stationary ARIMA(p, d, q) demand models. Under this framework, he shows that the propagation of demand is state-space-in, state-space-out; i.e., a given player's demand and order both have state-space representations. This forecasting approach can be applied to demand models that in the presence of an infinite past would not be invertible. Furthermore, using such an approach, one could determine the value of information sharing by comparing a supply chain player's MSFE under a sharing arrangement to the MSFE without a sharing arrangement.

Similar to the research by Aviv, Chen and Lee (2009) also consider a supply chain with a single retailer that uses a Martingale Model of Forecast Evolution (MMFE) to forecast demand. They focus on minimizing the supply chain costs and discuss the use of information sharing and various supply chain coordination approaches including information sharing and order variability control. They show the utility of a retailer sharing its projected future orders as opposed to its demand or its demand forecasts in order to reduce the overall supply chain costs. Their paper is closely related to our research since we both consider the tradeoff between the retailer's and supplier's inventory related costs. We however focus on constructing the entire efficient frontier between the standard deviations of the retailer's and supplier's inventory as opposed to identifying a single point in this frontier in which cumulative supply chain costs are optimized.

Our study of the tradeoff between the standard deviations of the retailer's inventory and the supplier's inventory is also related to the research of Graves et al. (1998). In their paper, they study the tradeoff frontier of the retailer's production smoothing and the standard deviation of the retailer's inventory. They use the single-stage problem as a building block for the more general multistage supply chain problem.

Cui et al. (2015) studies a supply chain where a supplier receives orders that are an aggregate of two processes due to the retailer placing an order that is generated by a standard inventory policy coupled with order smoothing and a decision deviation process. This assumption is based upon empirical evidence, where a planner does not place orders strictly according to a standard inventory policy but rather deviate from it for a variety of reasons. They conclude that in such a context, sharing the demand processes is almost always valuable. Without the additional process, they cite Giloni et al. (2014) in pointing out that there is value to information sharing if and only if the retailer's order is non-invertible. Indeed, in our paper, determining whether or not the retailer's

order is invertible with respect to its shocks is a crucial component. However, in our paper, the concept of the “value of information sharing” is more subtle than in the existing literature. The existing literature compares the inventory-related cost of a fixed policy with and without sharing. And it follows from Giloni et al. (2014) that the two costs are the same if and only if the order is invertible with respect to the retailer’s shocks. In this paper, we seek to compare the costs (direct and indirect) of an optimal policy under sharing with a (possibly different) optimal policy under no sharing.

1.2. Notation and Definitions

Let $X = \{X_t : t \in \mathbb{Z}\}$ denote a discrete-time Gaussian and stationary stochastic process.

- **LINEAR PAST** $\mathcal{M}_X(t)$: We denote by $\mathcal{M}_X(t) := \overline{\text{span}}\{1, X_t, X_{t-1}, \dots\}$ the linear past of X , that is, the Hilbert space generated by $\{1, X_t, X_{t-1}, \dots\}$ with inner product given by the covariance operator.

- **LINEAR FORECAST** $m_X(t)$: We define $m_X(t) := \mathbb{E}[X_{t+1} | \mathcal{M}_X(t)]$ to be the linear forecast of X_{t+1} conditional on the past history $\mathcal{M}_X(t)$.

- **MEAN SQUARED FORECAST ERROR (MSFE)** $\sigma_X^2(t)$:

$$\sigma_X^2(t) := \mathbb{E}[(X_{t+1} - m_X(t))^2 | \mathcal{M}_X(t)].$$

When X is a stationary process the MSFE is independent of t and we write σ_X^2 .

- **INVERTIBILITY**: We say that a process X is invertible with respect to a Gaussian white noise sequence $\xi = \{\xi_t : t \in \mathbb{Z}\}$ if $\mathcal{M}_X(t) = \mathcal{M}_\xi(t)$, that is, if the linear past of X coincides with the linear past of ξ .

- **CONVERGENCE EXPONENTIALLY FAST**: We say that a real-valued sequence $\{x_k\}_{k \geq 0}$ tends to zero exponentially fast if there exist positive constants c_1 and c_2 such that for all $k \geq 0$, $|x_k| \leq c_1 \exp(-c_2 k)$.

- **ORDER IN PROBABILITY** $O_p(1)$: We say that X is $O_p(1)$ if for every $\epsilon > 0$ there exist an M and N such that

$$\mathbb{P}(|X_t| > M) < \epsilon, \quad \text{for all } t > N.$$

2. Single-Stage Inventory Management System

In this section, we propose a mathematical model to establish the connection between the retailer’s ordering policy O_t and the variability of her on-hand inventory σ_I and the root mean squared forecast

error σ_S in the context of a single-stage inventory management system as depicted in Figure 4. The analysis and results of this single-stage model constitute a building block that we will use later, in Section 6, to study the two-tier supply chain system in Figure 1.




Figure 4 A single-state inventory system.

The retailer's demand is given by an exogenous stochastic process $\{D_t : t \in \mathbb{Z}\}$, where D_t is the market demand in period t . Throughout the paper, we will make the following assumption about this demand process.

Assumption 1 (MARKET DEMAND) *There exists a sequence $\Psi = \{\psi_n\}_{n \geq 0}$ with $\psi_0 = 1$ that tends to zero exponentially fast such that the retailer's demand $\{D_t\}$ admits a one-sided MA(∞) representation with respect to a Gaussian white noise sequence $\{\epsilon_t\}$ (with $\mathbb{E}[\epsilon_t] = 0$ and $\text{Var}[\epsilon_t] = \sigma_\epsilon^2$) given by*

$$D_t = d + \sum_{n=0}^{\infty} \psi_n \epsilon_{t-n}, \quad (1)$$

where $d > 0$ is the per-period mean demand. Furthermore, we assume that $\{D_t\}$ is invertible with respect to $\{\epsilon_t\}$.

Remark 1 A few comments about this assumption are in order. First, it is worth noticing that the representation of the demand process in equation (1) is not particularly restrictive as any stationary ARMA process admits such a representation. At the same time, the one-sided moving average representation does imply that $\{D_t\}$ is weakly stationary and so we are excluding non-stationary ARIMA processes and all other non-stationary processes. Second, under the stationarity assumption, the invertibility requirement entails no loss of generality since the demand to the retailer is exogenous so she sees only $\{D_t\}$ from which she must construct a set of shocks that span the same linear past as $\{D_t\}$, that is, $\mathcal{M}_D(t) = \mathcal{M}_\epsilon(t)$ (see Section 1.2 for definition). For the same reason, the normalization $\psi_0 = 1$ is also imposed without loss of generality to ensure the identifiability of the demand model. \square

The retailer serves the demand D_t from a stock of on-hand inventory that she replenishes by placing orders every period to an external supplier. We let O_t denote the order placed by the retailer in period t and assume a one-period replenishment lead-time, *i.e.*, the supplier delivers the order O_t in full at the beginning of period $t + 1$. We also assume that excess demand is fully

backlogged. It follows that the retailer's inventory I_t at the end of period t evolves according to the material-flow equation:

$$I_t = I_{t-1} + O_{t-1} - D_t \quad t = 1, 2, \dots \quad (2)$$

The retailer's objective is to minimize her long-term average procurement plus inventory costs. If we let $C_t^{\text{R}}(O_t, I_t)$ denote the retailer's cost in period t , as a function of her order quantity O_t and inventory level I_t , then her inventory management problem is given by:

$$\Pi^{\text{R}} := \inf_{O_t} \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T C_t^{\text{R}}(O_t, I_t) \right]. \quad (3)$$

In order to solve the optimization problem in (3) we need to further specify the structure of the retailer's cost function C_t^{R} . For example, it is a standard practice in the inventory management literature to assume a linear cost function of the form $C_t^{\text{R}}(O_t, I_t) = w O_t + h I_t^+ + b I_t^-$, where w is the per-unit procurement cost (or wholesale price) and h and b are the retailer's per-unit per-period holding and backordering costs, respectively. For the moment, instead of choosing a specific functional form for C_t^{R} , we will assume that the retailer's long-term average cost is a function of the volatility of the inventory σ_{I} and order process σ_{S} . (A precise definition of σ_{I} and σ_{S} will be given shortly). Mathematically, we assume that there exists an increasing function $\mathcal{C}^{\text{R}}(\sigma_{\text{I}}, \sigma_{\text{S}})$ such that

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T C_t^{\text{R}}(O_t, I_t) \right] = \mathcal{C}^{\text{R}}(\sigma_{\text{I}}, \sigma_{\text{S}}). \quad (4)$$

Implicit in this expression is the fact that both σ_{I} and σ_{S} depend on the retailer's inventory replenishment strategy $\{O_t\}$.

Our motivation to write the retailer's problem in (4) in terms of an abstract cost function $\mathcal{C}^{\text{R}}(\sigma_{\text{I}}, \sigma_{\text{S}})$ is to both highlight and isolate the two distinctive cost effects associated with an inventory replenishment policy that we discussed in the Introduction: (i) the *direct* operating costs associated with the retailer's on-hand inventory and (ii) the *indirect* replenishment costs associated with the variability of her orders. Later, in Section 6, we will consider a concrete example to illustrate how these two effects can be quantified and how one can derive equation (4) in the case that the retailer's cost takes the standard linear form $C_t^{\text{R}}(O_t, I_t) = w O_t + h I_t^+ + b I_t^-$. We also refer the reader to [Bray and Mendelson \(2012\)](#) for another example of a two-tier supply chain model in which the retailer optimizes a cost function that admits a representation like the one in equation (4) in which the functional \mathcal{C}^{R} is linear in σ_{I} and σ_{S} .

2.1. Inventory Variability: σ_1^2

In this section, we provide an expression for σ_1 in terms of the demand and order process. We start by making the following assumption:

Assumption 2 (RETAILER'S INVENTORY REPLENISHMENT POLICY) *The retailer's orders $\{O_t\}$ admit a one-sided MA(∞) representation*

$$O_t = \tilde{d} + \sum_{n=0}^{\infty} \tilde{\psi}_n \epsilon_{t-n}, \quad (5)$$

with respect to the white noise sequence $\{\epsilon_t\}$ defined in Assumption 1 and where the sequence $\tilde{\Psi} = \{\tilde{\psi}_n\}_{n \geq 0}$ tends to zero exponentially fast.

Remark 2 Under Assumption 1, the demand process is invertible and this guarantees that the retailer is able to infer her demand shocks $\{\epsilon_t\}$ and therefore select an order process as in equation (5). Note that the assumption implies that O_t is a stationary process. It is also worth pointing out that since O_t is based on the same shocks observed by the retailer, that is, O_t contains no additional exogenous source of variation introduced by the retailer (see Cui et al. (2015) for a model in which the retailer randomizes her orders and introduces additional random shocks). \square

To ensure that the retailer's inventory process in equation (2) is stable under a replenishment policy defined by (5), we have to restrict the set of values of \tilde{d} and $\{\tilde{\psi}_n\}_{n \geq 0}$ over which the retailer's minimizes her average long-term cost.

Definition 1 (Admissible Policy) *A replenishment policy $(\tilde{d}, \{\tilde{\psi}_n\}_{n \geq 0})$ is said to be admissible if $\{\tilde{\psi}_n\}_{n \geq 0}$ tends to zero exponentially fast and the following two conditions are satisfied: (i) $\tilde{d} = d$ and (ii) $\sum_{n=0}^{\infty} \tilde{\psi}_n = \sum_{n=0}^{\infty} \psi_n$.*

Condition (i) simply requires that the mean amount ordered per period \tilde{d} must be equal to the mean demand per period d . This is a natural requirement since otherwise the inventory would either grow or decrease unboundedly depending on whether $\tilde{d} > d$ or $\tilde{d} < d$, respectively. For this reason, in what follows we set $\tilde{d} = d$ and drop the dependence of the retailer's ordering strategy on \tilde{d} and refer to $\tilde{\Psi} = \{\tilde{\psi}_n\}_{n \geq 0}$ as the retailer's inventory replenishment policy.

Condition (ii) is also needed to ensure that the variance of the inventory process I_t remains bounded. To see this, first note that by iterating the inventory dynamics in equation (2) we get that I_t satisfies

$$I_t = I_0 + \sum_{n=1}^t (O_{n-1} - D_n) = I_0 + D_0 - D_t + \sum_{n=0}^{t-1} r_n, \quad (6)$$

where $r_n := O_n - D_n = \sum_{k=0}^{\infty} (\tilde{\psi}_k - \psi_k) \epsilon_{n-k}$.

Hence, for the variance of I_t to remain bounded we need the variance of the sum of the $\{r_n\}$ in equation (6) to remain bounded as t goes to infinity. To verify this condition we make use of the following intermediate result.

Lemma 1 *For a fixed admissible ordering policy, $\tilde{\Psi} = \{\tilde{\psi}_n\}_{n \geq 0}$, define the sequence $\{g_n\}_{n \geq 0}$ by*

$$g_0 := 0 \quad \text{and} \quad g_n := \sum_{k=0}^{n-1} (\tilde{\psi}_k - \psi_k), \quad n = 1, 2, \dots$$

Under the conditions in Assumptions 1 and 2

$$\sum_{n=0}^{t-1} r_n = \sum_{n=1}^{\infty} g_n \epsilon_{t-n} - \sum_{n=1}^{\infty} g_n \epsilon_{-n}.$$

Note that this expression is $O_p(1)$ since $g_L \rightarrow 0$ exponentially fast by Assumptions 1 and 2. The first term is a stationary time series, the second term is a constant and the covariance between the first and second terms goes to zero as $t \rightarrow \infty$. We summarize the previous discussion in the following proposition.

Proposition 1 *Under Assumptions 1 and 2, the variance of the inventory process I_t is bounded for any admissible replenishment policy $\tilde{\Psi} = \{\tilde{\psi}_n\}_{n \geq 0}$.*

Remark 3 (Initial Inventory) Combining the representation of the retailer's inventory process in equation (6) and the result in Lemma 1, we get the following alternative representation for I_t ,

$$I_t = I_0 + D_0 - \sum_{n=1}^{\infty} g_n \epsilon_{-n} + \sum_{n=1}^{\infty} g_n \epsilon_{t-n} - D_t = \tilde{I}_0 + \sum_{n=1}^{\infty} g_n \epsilon_{t-n} - (D_t - d), \quad (7)$$

where $\tilde{I}_0 := I_0 + D_0 - d - \sum_{n=1}^{\infty} g_n \epsilon_{-n}$.

From Remark 1, it follows that $\tilde{I}_0 \in \mathcal{M}_D(0)$. Hence, we can view \tilde{I}_0 as an additional degree of freedom that the retailer can set to further minimize her inventory cost. We will return to this point in Section 6.2. \square

From the MA(∞) representation of D_t in equation (1) and the representation of I_t in Remark 3, we get that

$$I_t = \tilde{I}_0 + \sum_{n=0}^{\infty} (g_n - \psi_n) \epsilon_{t-n}.$$

Hence, conditional on $\mathcal{M}_D(0)$, the retailer's inventory I_t is normally distributed with mean and variance given by

$$\mathbb{E}[I_t|\mathcal{M}_D(0)] = \tilde{I}_0 + \sum_{n=t}^{\infty} (g_n - \psi_n) \epsilon_{t-n} \quad \text{and} \quad \text{Var}(I_t|\mathcal{M}_D(0)) = \sigma_\epsilon^2 \sum_{n=0}^{t-1} (g_n - \psi_n)^2,$$

respectively. It follows that, conditional on $\mathcal{M}_D(0)$, I_t converges in distribution to a Normal random variable I_∞ with mean and variance

$$\mu_1 := \lim_{t \rightarrow \infty} \mathbb{E}[I_t|\mathcal{M}_D(0)] = \tilde{I}_0 \quad \text{and} \quad \sigma_1^2 := \lim_{t \rightarrow \infty} \text{Var}(I_t|\mathcal{M}_D(0)) = \sigma_\epsilon^2 \sum_{n=0}^{\infty} (g_n - \psi_n)^2, \quad (8)$$

respectively. When convenient, we will write $\sigma_1^2(\tilde{\Psi})$ to emphasize the dependence of the inventory volatility on the replenishment policy $\tilde{\Psi} = \{\tilde{\psi}_n\}_{n \geq 0}$.

2.2. Mean Squared Forecast Error: σ_S^2

Let us now turn to the derivation of $\sigma_S^2(\tilde{\Psi})$, the mean squared forecast error associated with the replenishment policy $\tilde{\Psi} = \{\tilde{\psi}_n\}_{n \geq 0}$. To this end, let $\mathcal{M}_S(t)$ be the supplier's information available in period t . Given this information, the supplier's best linear forecast of the retailer's order O_t is equal to $m_S(t) := \mathbb{E}[O_{t+1}|\mathcal{M}_S(t)]$ and the mean squared forecast error (MSFE) is given by

$$\sigma_S^2 = \text{Var}[(O_{t+1} - m_S(t))|\mathcal{M}_S(t)] \quad t = 1, 2, \dots \quad (9)$$

(The fact that σ_S^2 is independent of t follows from the stationarity of O_t in Assumption 2.)

In general, the information contained in $\mathcal{M}_S(t)$ is at least as large as the information generated by the linear past of the retailer's orders $\{O_t\}$ and at the same time it is contained in the information generated by the retailer's demand shocks $\{\epsilon_t\}$, that is, $\mathcal{M}_O(t) \subseteq \mathcal{M}_S(t) \subseteq \mathcal{M}_\epsilon(t)$. For the purpose of this paper, we will only consider the two extreme cases (see [Kovtun et al. \(2014\)](#) for further discussion on other cases with intermediate levels of information) :

- **FULL INFORMATION SHARING:** $\mathcal{M}_S(t) = \mathcal{M}_\epsilon(t)$. In this case, the retailer is willing and able to share her demand shock information with the supplier. It follows from (5) that the supplier's linear forecast is equal to $m_S(t) = d + \sum_{n=1}^{\infty} \tilde{\psi}_n \epsilon_{t+1-n}$ and so the mean squared forecast error in equation (9) reduces to

$$\sigma_{S|FI}^2(\tilde{\Psi}) := \tilde{\psi}_0^2 \sigma_\epsilon^2. \quad (10)$$

(The subscript 'FI' is mnemonic for Full Information.)

• **NO INFORMATION SHARING:** $\mathcal{M}_s(t) = \mathcal{M}_o(t)$. In this case, the only information available to the supplier is the one generated by the retailer's orders $\{O_t\}$. We can use Kolmogorov's formula (see §5.8 in Brockwell and Davis 2006) to express σ_s^2 in equation (9) in terms of the retailer's replenishment policy parameters $\tilde{\Psi} = \{\tilde{\psi}_k\}_{k \geq 0}$ to get

$$\sigma_{s|NS}^2(\tilde{\Psi}) := \sigma_\epsilon^2 \exp\left(\frac{1}{2\pi} \int_{-\pi}^{\pi} \log\left(\tilde{\psi}(e^{-i\lambda})\tilde{\psi}(e^{i\lambda})\right) d\lambda\right), \quad \text{where } \tilde{\psi}(z) := \sum_{n=0}^{\infty} \tilde{\psi}_n z^n \quad (11)$$

is the autocovariance generating function of the retailer's orders $\{O_t\}$. (The subscript 'NS' is mnemonic for No Information Sharing.)

Using the results in (10) and (11), we can derive the following simple bounds for the value of the supplier's mean squared forecast error under no information sharing as a function of the replenishment policy parameters $\tilde{\Psi} = \{\tilde{\psi}_n\}_{n \geq 0}$.

Lemma 2 (Bounds on MSFE) *The mean squared forecast errors $\sigma_{s|FI}^2(\tilde{\Psi})$ and $\sigma_{s|NS}^2(\tilde{\Psi})$ satisfy*

$$\tilde{\psi}_0^2 \sigma_\epsilon^2 = \sigma_{s|FI}^2(\tilde{\Psi}) \leq \sigma_{s|NS}^2(\tilde{\Psi}) \leq \sigma_\epsilon^2 \sum_{n=0}^{\infty} \tilde{\psi}_n^2.$$

Example 1 *Consider the case in which the retailer's orders are generated according to the MA(1) process $O_t = d + \epsilon_t + \theta \epsilon_{t-1}$, for some parameter $\theta \in \mathbb{R}$ and assume there is no information sharing. Figure 5 depicts the value of the supplier's mean squared forecast error σ_s^2 as well as the bounds in Lemma 2 as functions of θ for the case in which the innovation variance is equal to one.*

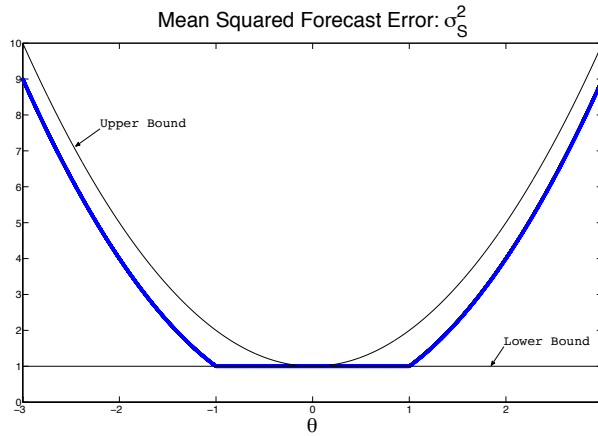


Figure 5 Supplier's mean squared forecast error under no information sharing when orders follows an MA(1) process $O_t = d + \epsilon_t + \theta \epsilon_{t-1}$ with $\sigma_\epsilon = 1$.

A few remarks about this figure are in order. First, it is worth noticing that when $\theta \in [-1, 1]$, the lower bound is tight. This is due to the fact that for $\theta \in [-1, 1]$ the corresponding order process

$\{O_t\}$ is invertible with respect to the demand shocks $\{\epsilon_t\}$ and so $\mathcal{M}_O(t) = \mathcal{M}_\epsilon(t)$. In other words, information sharing has no impact on the supplier's available information when $\theta \in [-1, 1]$ as he can recover the demand shocks $\{\epsilon_t\}$ directly by inverting the retailer's order process. It follows that $\sigma_{S|NS}^2 = \sigma_{S|FI}^2$ in this case.

On the flip side, for $\theta \notin [-1, 1]$, the order process is not invertible and the value of $\sigma_{S|NS}^2$ gets closer to the upper bound as θ increases. Indeed, using the expression for $\sigma_{S|NS}^2$ in equation (11), one can show that

$$\lim_{|\theta| \rightarrow \infty} \frac{\sigma_\epsilon^2 \sum_{n=0}^{\infty} \tilde{\psi}_n^2}{\sigma_{S|NS}^2} = 1. \quad \square$$

3. Inventory Policies under Full Information Sharing

Now that we have expressed the values of $\sigma_I(\tilde{\Psi})$ and $\sigma_S(\tilde{\Psi})$ in terms of the retailer's replenishment policy $\tilde{\Psi} = \{\tilde{\psi}_n\}_{n \geq 0}$, we can turn to the question of how the retailer should select the parameters $\{\tilde{\psi}_n\}$ to minimize her average cost. To answer this question we would need to specify the cost function $\mathcal{C}^R(\sigma_I, \sigma_S)$ appearing in equation (4). Instead, we will tackle this problem in a different way by characterizing a Pareto *efficient frontier* \mathcal{E} of pairs $\{(\sigma_I^*, \sigma_S^*)\}$ such that for any increasing cost function \mathcal{C}^R the retailer's optimal ordering strategy is such that the corresponding root mean squared forecast error σ_S and inventory volatility σ_I belong to this Pareto frontier. As we will see, our ability to compute \mathcal{E} depends critically on the type of collaboration between the retailer and the supplier. If the retailer is willing to share her demand information then a simple representation of the Pareto frontier can be obtained. On the other hand, under no information sharing the problem is significantly more involved and we are only able to provide a partial characterization of \mathcal{E} .

To this end, let us consider the following optimization problem:

$$\Sigma(\sigma) := \inf_{\tilde{\Psi} = \{\tilde{\psi}_n\}_{n \geq 0}} \left\{ \sigma_I(\tilde{\Psi}) \text{ subject to } \sigma_S(\tilde{\Psi}) \leq \sigma \text{ and } \sum_{n=0}^{\infty} \tilde{\psi}_n = \sum_{n=0}^{\infty} \psi_n \right\}.$$

In words, this optimization searches (over the set of admissible ordering policies $\{\tilde{\psi}_n\}_{n \geq 0}$) for the minimum value of σ_I subject to the constraint that σ_S must be bounded above by the value σ . Since we are assuming that the cost function $\mathcal{C}^R(\sigma_I, \sigma_S)$ is increasing in both arguments, an optimal replenishment policy $\tilde{\Psi}^*$ must satisfy $\sigma_I(\tilde{\Psi}^*) = \Sigma(\sigma_S(\tilde{\Psi}^*))$.

It follows from the discussion in the previous section that the complexity of characterizing the mapping $\Sigma(\cdot)$ depends directly on the type of information available to the supplier. Let us first discuss the more tractable case with full information sharing, $\mathcal{M}_S(t) = \mathcal{M}_\epsilon(t)$, in this section and postpone to the following section the case in which there is no information sharing.

Using the representations of σ_I and $\sigma_{S|FI}$ in equations (8) and (10), respectively, we can rewrite the optimization problem that defines $\Sigma(\cdot)$ under full information sharing as follows:

$$\Sigma_{FI}(\sigma) := \inf_{\tilde{\Psi}} \sigma_\epsilon \sqrt{\psi_0^2 + \sum_{n=1}^{\infty} (g_n - \psi_n)^2} \quad (12)$$

$$\text{subject to } \sigma_\epsilon |\tilde{\psi}_0| \leq \sigma, \quad (13)$$

$$\sum_{n=0}^{\infty} \tilde{\psi}_n = \sum_{n=0}^{\infty} \psi_n, \quad g_0 = 0 \quad \text{and} \quad g_k = \sum_{n=0}^{k-1} (\tilde{\psi}_n - \psi_n), \quad k = 1, 2, \dots \quad (14)$$

Although this optimization has an infinite number of decision variables, it admits a surprisingly simple solution.

Proposition 2 *Under full information, when $\mathcal{M}_S(t) = \mathcal{M}_\epsilon(t)$,*

$$\Sigma_{FI}(\sigma) = \sigma_\epsilon \sqrt{\psi_0^2 + \left(\frac{\sigma}{\sigma_\epsilon} - |\psi_0 + \psi_1| \right)^2}$$

and the corresponding optimal solution $\tilde{\Psi} = \{\tilde{\psi}_n\}_{n \geq 0}$ is given by

$$\tilde{\psi}_0 = \text{sign}(\psi_0 + \psi_1) \frac{\sigma}{\sigma_\epsilon}, \quad \tilde{\psi}_1 = \psi_2 + \psi_1 + \psi_0 - \tilde{\psi}_0 \quad \text{and} \quad \tilde{\psi}_n = \psi_{n+1} \quad \text{for all } n \geq 2. \quad (15)$$

Note that the optimal replenishment policy in (15) is parametrized by the value of $\sigma \in \mathbb{R}_+$. By varying this parameter, we can define the set \mathcal{E}_{FI} of efficient pairs (σ_S, σ_I) in the sense of Pareto optimality. That is, since the retailer would like to simultaneously minimize both types of variabilities, there is a subset of volatility pairs (σ_S, σ_I) that are efficient in the sense that they collectively dominate every other achievable pair. Using the result in Proposition 2, we characterize this efficient frontier in the following Definition.

Definition 2 (Efficient Replenishment under Full Information) *Given a demand process D_t with $MA(\infty)$ representation as in equation (1), we say that an inventory replenishment policy O_t is efficient if the resulting mean squared forecast error and inventory variance pair $(\sigma_{S|FI}^2, \sigma_I^2)$ induced by O_t belongs to the set*

$$\mathcal{E}_{FI} := \left\{ (\sigma_S, \sigma_I) : 0 \leq \sigma_S \leq \sigma_\epsilon |\psi_0 + \psi_1| \quad \text{and} \quad \sigma_I = \sigma_\epsilon \sqrt{\psi_0^2 + \left(\frac{\sigma_S}{\sigma_\epsilon} - |\psi_0 + \psi_1| \right)^2} \right\}.$$

Interestingly, the set \mathcal{E}_{FI} depends on the parameters of the demand process $\Psi = \{\psi_n\}_{n \geq 0}$ only through the first two terms ψ_0 and ψ_1 and is independent of the values $\{\psi_n\}_{n \geq 2}$. An example of \mathcal{E}_{FI} for the case in which market demand is IID (with $\sigma_\epsilon = \psi_0 = 1$ and $\psi_n = 0$ for all $n \geq 1$) was depicted in Figure 2.

Extreme Cases: There are two distinctive efficient inventory replenishment policies that are worth highlighting. They correspond to the extreme points of the efficient frontier \mathcal{E}_{FI} .

-) PREDICTABLE ORDERS: (P) The retailer can choose a replenishment policy that completely eliminates the supplier's root mean squared forecast error making her orders completely predictable. From the result in Proposition 2, this can be achieved by setting $\tilde{\psi}_0^{\text{P}} = 0$, $\tilde{\psi}_1^{\text{P}} = \psi_2 + \psi_1 + \psi_0$ and $\tilde{\psi}_k^{\text{P}} = \psi_{k+1}$ for all $k \geq 2$. (The superscript 'P' is mnemonic for Predictable.) Under this Predictable policy the root mean squared forecast error and inventory volatility are equal to

$$\sigma_{\text{S|FI}}^{\text{P}} = 0 \quad \text{and} \quad \sigma_{\text{I|FI}}^{\text{P}} = \sqrt{\psi_0^2 + (\psi_0 + \psi_1)^2} \sigma_\epsilon.$$

Note that this predictable strategy corresponds to point 'P' in Figure 2. Also, under the choice $\tilde{\psi}_0^{\text{P}} = 0$, we have $\mathcal{M}_O(t) \subsetneq \mathcal{M}_\epsilon(t)$ and so the retailer must engage in information sharing to ensure the full information requirement and the invertibility of the order process.

-) MYOPIC POLICY: (M) Suppose the retailer is only concerned with minimizing the variance of her inventory process. In this case, she would like to choose the replenishment policy corresponding to point 'M' in Figure 2. According to Proposition 2, this can be achieved by choosing $\tilde{\psi}_0^{\text{M}} = \psi_0 + \psi_1$ and $\tilde{\psi}_n^{\text{M}} = \psi_{n+1}$ for all $n \geq 1$. (In this case, the superscript 'M' is mnemonic for Myopic.) Under this Myopic policy the root mean squared forecast error and inventory volatility are equal to

$$\sigma_{\text{S|FI}}^{\text{M}} = |\tilde{\psi}_0^{\text{M}}| \sigma_\epsilon = |\psi_0 + \psi_1| \sigma_\epsilon \quad \text{and} \quad \sigma_{\text{I|FI}}^{\text{M}} = |\psi_0| \sigma_\epsilon.$$

Remark 4 We call this replenishment policy 'Myopic' because it coincides with the classical myopic order-up-to policy (see Johnson and Thompson 1975 for details[†]). We will elaborate further on this connection in Section 6. \square

4. Inventory Policies with No Information Sharing

We turn to the case in which the retailer is not able (or willing) to share her shocks with the supplier. Using the expression for $\sigma_{\text{S|NS}}$ in equation (11), we can formulate the optimization problem that defines $\Sigma(\sigma)$ under no information sharing as follows:

$$\Sigma_{\text{NS}}(\sigma) := \inf_{\tilde{\Psi}} \sigma_\epsilon \sqrt{\psi_0^2 + \sum_{n=1}^{\infty} (g_n - \psi_n)^2} \quad (16)$$

[†] Myopic policies have a long tradition in the inventory management literature (e.g., A.F. Veinott 1965) as they offer a simple and tractable approximation to what is, otherwise, a challenging optimization problem to solve in general. Furthermore, an important part of their appeal is that they can be optimal under suitable conditions on the underlying demand process.

$$\text{subject to } \sigma_\epsilon^2 \exp\left(\frac{1}{2\pi} \int_{-\pi}^{\pi} \log\left(\tilde{\psi}(e^{-i\lambda})\tilde{\psi}(e^{i\lambda})\right) d\lambda\right) \leq \sigma^2, \quad (17)$$

$$\sum_{n=0}^{\infty} \tilde{\psi}_n = \sum_{n=0}^{\infty} \psi_n, \quad g_0 = 0 \quad \text{and} \quad g_k = \sum_{n=0}^{k-1} (\tilde{\psi}_n - \psi_n), \quad k = 1, 2, \dots \quad (18)$$

Contrary to the case with full information sharing, a general solution to this infinite-dimensional optimization problem has proven elusive. Its complexity lies essentially in the first constraint (17) that connects $\tilde{\Psi} = \{\tilde{\psi}_n\}_{n \geq 0}$ to the mean squared forecast error $\sigma_{\text{S|NS}}^2$. There are, however, situations in which a solution can be obtained using the methods developed in the previous section and we will focus on these cases first.

By Lemma 2 we know that for any $\tilde{\Psi} = \{\tilde{\psi}_n\}_{n \geq 0}$, we have $\sigma_{\text{S|FI}}^2(\tilde{\Psi}) \leq \sigma_{\text{S|NS}}^2(\tilde{\Psi})$. As a result, the left-hand side in constraint (17) is bounded below by $\sigma_\epsilon^2 \tilde{\psi}_0^2$. This means that any feasible solution to (16)-(18) is also a feasible solution to the optimization problem (12)-(14) under full information sharing. The following result is a direct consequence of this observation.

Lemma 3 $\Sigma_{\text{NS}}(\sigma) \geq \Sigma_{\text{FI}}(\sigma)$, for all $\sigma \in \mathbb{R}_+$.

This lemma implies that the efficient frontier \mathcal{E}_{FI} under full information sharing in Definition 2 is a lower bound to the corresponding frontier \mathcal{E}_{NS} under no information sharing. Now, consider a replenishment policy $\tilde{\Psi} = \{\tilde{\psi}_n\}_{n \geq 0}$ that is invertible in the sense that the order process $\{O_t\}$ induced by $\tilde{\Psi}$ is invertible with respect to the shocks $\{\epsilon_t\}$ (see Section 1.2). It follows then that $\sigma_{\text{S|FI}}^2(\tilde{\Psi}) = \sigma_{\text{S|NS}}^2(\tilde{\Psi})$ and the left-hand side in constraint (17) reduces to

$$\sigma_\epsilon^2 \exp\left(\frac{1}{2\pi} \int_{-\pi}^{\pi} \log\left(\tilde{\psi}(e^{-i\lambda})\tilde{\psi}(e^{i\lambda})\right) d\lambda\right) = \sigma_\epsilon^2 \tilde{\psi}_0^2.$$

This condition, together with Lemma 3 lead to the following result.

Proposition 3 Let $\tilde{\Psi} = \{\tilde{\psi}_n\}_{n \geq 0}$ be the optimal replenishment policy under full information derived in Proposition 2, that is,

$$\tilde{\psi}_0 = \text{sign}(\psi_0 + \psi_1) \frac{\sigma}{\sigma_\epsilon}, \quad \tilde{\psi}_1 = \psi_2 + \psi_1 + \psi_0 - \tilde{\psi}_0 \quad \text{and} \quad \tilde{\psi}_n = \psi_{n+1} \quad \text{for all } n \geq 2$$

for some fixed value $\sigma \geq 0$, and suppose that the order process generated by $\tilde{\Psi}$ is invertible with respect to the shocks $\{\epsilon_t\}$. Then, $\tilde{\Psi} = \{\tilde{\psi}_n\}_{n \geq 0}$ solves the optimization problem (16)-(18) and the solution coincides with the solution under full information sharing.

Remark 5 An immediate and important corollary of the previous result is that there is no real value in information sharing if the solution under full information sharing is invertible. The reason is that invertibility implies that $\mathcal{M}_O(t) = \mathcal{M}_\epsilon(t)$, that is, the supplier can recover all the relevant information known by the retailer by simply observing the history of orders $\{O_t\}$. \square

Another implication of Proposition 3 is that the efficient frontiers of volatilities pairs (σ_s, σ_1) under full information sharing and no information sharing coincide on those values of σ for which the order process is invertible with respect to the shocks $\{\epsilon_t\}$. Unfortunately, the question of how to identify those cases in which the optimal replenishment policy under sharing is invertible must be conducted on a case by case basis and boils down to verifying if the identify $\sigma_{s|FI}^2(\tilde{\Psi}) = \sigma_{s|NS}^2(\tilde{\Psi})$ holds or not. The following example illustrate this point.

Example 2 (IID Demand) *Let us consider the case in which the retailer's demand is white noise. That is, $\Psi = \{\psi_n\}_{n \geq 0}$ with $\psi_0 = 1$ and $\psi_n = 0$ for all $n \geq 1$. We note that this is an important special case since it (together with Gaussianity) implies that the retailer's per-period demands $\{D_t\}$ are iid, which is a predominant modeling assumption in the inventory management literature (e.g., Gavirneni et al. 1999). Also, for notational simplicity let us assume that $\sigma_\epsilon = 1$.*

In this case, the optimal inventory replenishment policy $\tilde{\Psi} = \{\tilde{\psi}_n\}_{n \geq 0}$ under full information sharing in Proposition 2 is given by an MA(1) process with

$$\tilde{\psi}_0 = \sigma_s, \quad \tilde{\psi}_1 = 1 - \sigma_s \quad \text{and} \quad \tilde{\psi}_n = 0 \quad \text{for all } n \geq 2,$$

where $0 \leq \sigma_s^2 \leq 1$ is the corresponding MSFE induced by $\tilde{\Psi}$. (We exclude values of σ_s greater than 1 since they lead to replenishment policies that are not in the efficient frontier, see Definition 2.)

Now, it is well known that the MA(1) process above is invertible if and only if $\sigma_s \geq 0.5$. Hence, by Proposition 3 we know that this MA(1) solution is also optimal for the case with no information sharing if $\sigma_s \in [0.5, 1]$. In other words, if the retailer would like to select an optimal replenishment policy with a root MSFE greater than 0.5 then she can select the MA(1) process above even if she is not able (or willing) to share any demand information with the supplier. \square

The previous example raises the question of how to identify an optimal (Pareto efficient) replenishment policy under no information sharing when the invertibility condition in Proposition 3 does not hold. At this point, we do not have a definitive analytical answer to this question and it appears to us that such an answer would have to be found using alternative methods to the one developed in this paper. Specifically, it is our impression that the optimization problem (16)-(18)

that uses the MA(∞) coefficients $\{\tilde{\psi}_n\}_{n \geq 0}$ to represent a replenishment policy might not be the most adequate representation of the problem. The issue with this parametrization is that it does not lead a tractable mathematical framework to characterize invertible processes. An alternative, possibly more promising parametrization of the problem, is to work directly with the roots of the characteristic polynomial associated with the retailer's replenishment policy since the issue of checking invertibility becomes trivial in this case. However, it is not clear to us whether this alternative approach would in general simplify the optimization problem (16)-(18) and our ability to find an optimal solution.

Instead of pursuing a different parametrization of the problem, in what follows we propose an approximate formulation that replaces the optimization problem (16)-(18) by one that is amenable to analysis and conventional optimization techniques.

The approach that we take is to work with an alternative representation of $\sigma_{S|NS}^2$ in terms of the MA coefficients $\{\tilde{\psi}_n\}$ than the one available using Kolmogorov's formula in equation (11). The basic idea is to work directly with the representation of $\sigma_{S|NS}^2$ as the minimum forecast error among linear forecasts. To this end, given a history of orders $\{O_k\}_{k \leq t}$, we construct the linear forecast

$$\hat{O}_{t+1} = d + \sum_{k=1}^{\infty} \alpha_k (O_{t+1-k} - d),$$

for some fixed vector of parameters $\alpha = (\alpha_n)_{n \geq 1}$ independent of t .

The squared error $(O_{t+1} - \hat{O}_{t+1})^2$ can be written in terms of the MA coefficients $\{\tilde{\psi}_n\}$ as follows

$$(O_{t+1} - \hat{O}_{t+1})^2 = \left(\sum_{n=0}^{\infty} \left[\tilde{\psi}_n - \sum_{k=0}^{n-1} \alpha_{n-k} \tilde{\psi}_k \right] \epsilon_{t+1-n} \right)^2.$$

Then, under appropriate summability requirements, the expected mean squared forecast error, as a function of the vector α , is equal to

$$\mathbb{E}[(O_{t+1} - \hat{O}_{t+1})^2] = \sigma_{\epsilon}^2 \sum_{n=0}^{\infty} \left[\tilde{\psi}_n - \sum_{k=0}^{n-1} \alpha_{n-k} \tilde{\psi}_k \right]^2.$$

When can then can compute the value of $\sigma_{S|NS}^2$ as follows:

$$\sigma_{S|NS}^2 = \inf_{\alpha} \sigma_{\epsilon}^2 \sum_{n=0}^{\infty} \left[\tilde{\psi}_n - \sum_{k=0}^{n-1} \alpha_{n-k} \tilde{\psi}_k \right]^2.$$

In practice, by truncating the vector α we can get an approximation for $\sigma_{S|NS}^2$ solving

$$\sigma_{S|NS}^2(N) := \min_{\alpha} \sigma_{\epsilon}^2 \sum_{n=0}^{\infty} \left[\tilde{\psi}_n - \sum_{k=0}^{n-1} \alpha_{n-k} \tilde{\psi}_k \right]^2 \quad \text{subject to} \quad \alpha_n = 0 \quad \forall n > N.$$

It is well-known (see Chapter 5, Corollary 5.1.1, in Brockwell and Davis 2006) that the solution to this optimization problem is equal to

$$\sigma_{\text{S|NS}}^2(N) = \sigma_\epsilon^2 (\gamma_0 - \Gamma_N M_N^{-1} \Gamma_N^T), \quad (19)$$

where

$$\gamma_k = \sum_{n=0}^{\infty} \tilde{\psi}_n \tilde{\psi}_{n+k}, \quad \Gamma_N = (\gamma_1, \gamma_2, \dots, \gamma_N) \quad \text{and} \quad M_N = [m_{ij}] \in \mathbb{R}^N \times \mathbb{R}^N \quad \text{with} \quad m_{ij} := \gamma_{|i-j|}.$$

Using the approximation in equation (19), we can replace constraint (17) and approximate the optimization problem (16)-(18) by

$$\begin{aligned} \Sigma_{\text{NS}}(\sigma) &\approx \inf_{\tilde{\Psi}} \sigma_\epsilon \sqrt{\psi_0^2 + \sum_{n=1}^{\infty} (g_n - \psi_n)^2} \\ \text{subject to} \quad &\sigma_\epsilon^2 (\gamma_0 - \Gamma_N M_N^{-1} \Gamma_N^T) \leq \sigma^2, \\ &\sum_{n=0}^{\infty} \tilde{\psi}_n = \sum_{n=0}^{\infty} \psi_n, \quad g_0 = 0 \quad \text{and} \quad g_k = \sum_{n=0}^{k-1} (\tilde{\psi}_n - \psi_n), \quad k = 1, 2, \dots \end{aligned}$$

Based on a number of numerical experiments, solving the previous problem using standard optimization packages can be unstable in some instances. The more direct approach of solving for both α and $\tilde{\psi}$ simultaneously has shown a much better numerical performance. For this reason, we propose the following alternative formulation:

$$\Sigma_{\text{NS}}^2(\sigma) \approx \inf_{\tilde{\Psi}, \alpha} \left\{ \sigma_\epsilon^2 \sum_{n=0}^{\infty} \left(\sum_{k=0}^n \psi_n - \sum_{k=0}^{n-1} \tilde{\psi}_k \right)^2 \right\} \quad (20)$$

$$\text{subject to} \quad \sigma_\epsilon^2 \sum_{n=0}^{\infty} \left(\tilde{\psi}_n - \sum_{k=0}^{n-1} \alpha_{n-k} \tilde{\psi}_k \right)^2 \leq \sigma^2 \quad \text{and} \quad \sum_{n=0}^{\infty} \tilde{\psi}_n = \sum_{k=0}^{\infty} \psi_n. \quad (21)$$

As stated, problem (20)-(21) is an infinite dimensional optimization problem. In practice, we must truncate the sequences $\{\tilde{\psi}_n\}_{n \geq 0}$ and $\alpha = (\alpha_n)_{n \geq 1}$ to compute an approximate solution. For example, if we assume that demand is given by an MA(q) model for some $q \in \mathbb{N}$, we can solve this problem with the additional restrictions that orders are also given by an MA(Q) for some $Q \geq q$ (i.e., $\tilde{\psi}_n = 0$ for all $n > Q$) and that $\alpha_n = 0$ for all $n \geq N + 1$ for some $N \in \mathbb{N}$.

Example 3 (IID Demand - Continued) *Consider again the case in which the retailer's demand is white noise with $\psi_0 = 1$, $\psi_n = 0$ for all $n \geq 1$ and $\sigma_\epsilon = 1$, and let us compare the efficient frontiers*

\mathcal{E}_{FI} and \mathcal{E}_{NS} of volatility pairs (σ_S, σ_I) under full information sharing and no information sharing, respectively. From Definition 2, we have that

$$\mathcal{E}_{\text{FI}} := \left\{ (\sigma_S, \sigma_I) : 0 \leq \sigma_S \leq 1 \quad \text{and} \quad \sigma_I = \sqrt{1 + (\sigma_S - 1)^2} \right\}.$$

On the other hand, we use the optimization (20)-(21) to approximately compute \mathcal{E}_{NS} . The following figure depicts \mathcal{E}_{FI} and the approximation for \mathcal{E}_{NS} with $Q = 10$ and $N = 75$.

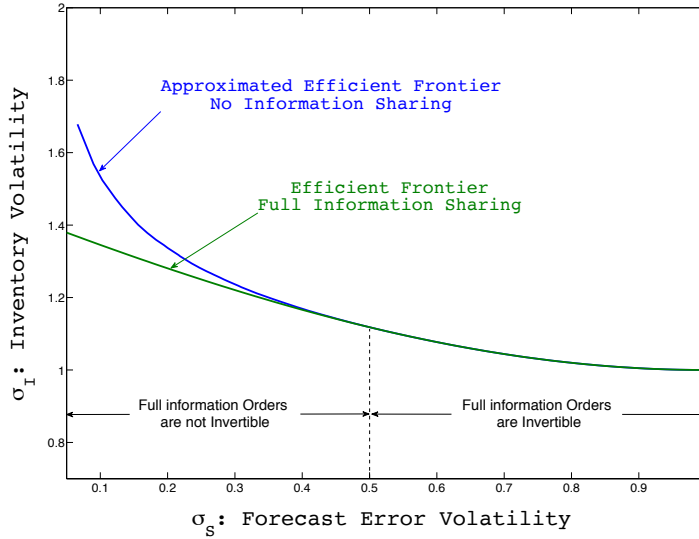


Figure 6 Efficient frontiers of volatility pairs (σ_S, σ_I) under full information sharing and no information sharing. Data $\sigma_\epsilon = \psi_0 = 1$ and $\psi_k = 0$ for all $k \geq 1$.

Recall that the optimal replenishment policy under full information sharing is invertible for values of $\sigma_S \geq 0.5$. In this region, the two frontiers coincide which is consistent with our previous discussion, i.e., for $\sigma_S \geq 0.5$ information sharing offers no additional value. On the other hand, in the region $\sigma_S < 0.5$ the efficient frontier under full information sharing is strictly lower and so information sharing is valuable in this case. Of course, we must keep in mind that the value of \mathcal{E}_{NS} depicted in Figure 6 is only an approximation (actually an upper bound) on the true value of the efficient frontier under no information sharing. \square

5. Numerical Experiments

To illustrate our methodology and results, in this section we conduct a set of numerical experiments to compare the set of efficient replenishment policies under full information sharing and no

information sharing to four alternative order smoothing policies that have been proposed in the literature.

Benchmark Policies: For completeness, we provide a brief mathematical description of these four benchmark policies and refer the reader to the original papers that proposed them for further details.

a) SIMPLE MOVING AVERAGE: (MA) Suppose the retailer's inventory strategy belongs to the family of simple moving average:

$$O_t^{\text{MA}} = \frac{1}{N} \sum_{n=0}^{N-1} D_{t-n}, \quad (22)$$

for some fixed integer $N \geq 1$. The use of a simple moving average ordering policy in the context of a two-tier supply chain system was studied by Balakrishnan, Geunes, and Pangburn (2004) for the special case in which market demand is IID. In our context, using the $\text{MA}(\infty)$ representation of D_t in equation (1) we get that

$$O_t^{\text{MA}} = d + \sum_{k=0}^{\infty} \tilde{\psi}_k^{\text{MA}} \epsilon_{t-k}, \quad \text{with} \quad \tilde{\psi}_k^{\text{MA}} := \frac{1}{N} \sum_{n=(k+1-N)^+}^k \psi_n.$$

It is not hard to see that O_t^{MA} satisfies the conditions in Definition 1 and so it is an admissible replenishment policy. Based on the representation of σ_I in equation (8) and after some tedious but straightforward manipulations we get that the root mean squared forecast error (under full information sharing) and inventory variance under a moving average policy are given by

$$\sigma_{\text{S|FI}}^{\text{MA}} = \frac{\sigma_\epsilon}{N} |\psi_0| \quad \text{and} \quad \sigma_I^{\text{MA}} = \frac{\sigma_\epsilon}{N} \sqrt{\sum_{k=0}^{\infty} \left(\sum_{n=(k+1-N)^+}^k (N-k+n) \psi_n \right)^2}.$$

In general, the pair $(\sigma_{\text{S|FI}}^{\text{MA}}, \sigma_I^{\text{MA}})$ is not efficient in the sense of Definition 2. Nevertheless, moving average policies have a practical appeal due to their simplicity. Also, we see from the expression for $\sigma_{\text{S|FI}}^{\text{MA}}$ that as N grows the mean forecast error decreases monotonically and the retailer's orders will become more predictable. One can also show that the moving average policy with $N = 1$ coincides with the myopic policy for the case in which $\{D_t\}$ is an IID sequence (*i.e.*, $\psi_k = 0$ for all $k \geq 1$).

b) EXPONENTIAL SMOOTHING: (ES) This is another order smoothing policy proposed by Balakrishnan et al. (2004). In this case, the retailer's replenishment policy is

$$O_t^{\text{ES}} = (1 - \rho) \sum_{n=0}^{\infty} \rho^n D_{t-n},$$

for some fixed $\rho \in [0, 1)$. From the MA(∞) representation of D_t in equation (1) we get that

$$O_t^{\text{ES}} = d + \sum_{k=0}^{\infty} \tilde{\psi}_k^{\text{ES}} \epsilon_{t-k}, \quad \text{with} \quad \tilde{\psi}_k^{\text{ES}} := (1 - \rho) \sum_{n=0}^k \rho^{k-n} \psi_n.$$

One can easily verify that O_t^{ES} is admissible in the sense of Definition 1. In this case, the mean forecast error and inventory volatilities are equal to

$$\sigma_{\text{S|FI}}^{\text{ES}} = \sigma_\epsilon (1 - \rho) |\psi_0| \quad \text{and} \quad \sigma_I^{\text{ES}} = \sigma_\epsilon \sqrt{\sum_{k=0}^{\infty} \left(\sum_{n=0}^k \rho^{k-n} \psi_n \right)^2}.$$

c) GKH: In a paper by [Graves, Kletler, and Hetzel \(1998\)](#), a production smoothing policy is proposed with the objective of minimizing the variance of the order process subject to a constraint that upper bounds the variance of the inventory process. The demand model considered by [Graves et al. \(1998\)](#) is more general than ours but a special case that coincides with ours is studied in detail in Section 2 of that paper.

The GKH policy is defined by a matrix of weights $W = [w_{ij}]$ such that

$$O_t^{\text{GKH}} = d + \sum_{k=0}^{\infty} \tilde{\psi}_k^{\text{GKH}} \epsilon_{t-k}, \quad \text{with} \quad \tilde{\psi}_k^{\text{GKH}} := \sum_{j=0}^H w_{kj} \psi_j,$$

where H is a non-negative integer that parametrizes how far into the future the retailer's updates her demand forecasts. The value of W are computed by solving the optimization problem

$$\min_{w_{ij}} \sum_{i=0}^H \sum_{j=0}^H (w_{ij} \psi_j \sigma_\epsilon)^2 \quad \text{subject to} \quad \text{Var}(I_t) \leq K^2 \quad \text{and} \quad \sum_{i=0}^H w_{ij} = 1 \quad \forall j,$$

where K is a parameter that bounds the variance of the retailer's inventory. [Graves et al. \(1998\)](#) provide a solution to this optimization.

d) CL: [Chen and Lee \(2009\)](#) studied a two-tier supply chain system similar to the one we consider and proposed an order smoothing policy that minimizes the cumulative supply chain (retailer and supplier) costs. On the one hand, the mathematical framework that they consider allows for more general demand models and replenishment lead-times. On the other hand, their analysis does not explicitly address the issue of invertibility.

In the context of our model, the CL policy is equal to

$$O_t^{\text{CL}} = d + \sum_{k=0}^{\infty} \tilde{\psi}_k^{\text{CL}} \epsilon_{t-k}, \quad \text{with} \quad \tilde{\psi}_0^{\text{CL}} := (1 - \gamma)(\psi_0 + \psi_1), \quad \tilde{\psi}_1^{\text{CL}} := \psi_2 + \gamma(\psi_0 + \psi_1), \quad \tilde{\psi}_n^{\text{CL}} := \psi_{n+1} \quad \forall n \geq 2,$$

where the parameter $\gamma \in [0, 1]$ is a function of the relative cost structures of the supplier and retailer as well as a relative measure of uncertainty faced by the supplier (see [Chen and Lee \(2009\)](#) for details).

Efficient Frontier: Let us now compare the values of (σ_s, σ_I) that are achievable under the four benchmark policies to the efficient frontiers with and without information sharing. In evaluating the performance of the four benchmark policies, we have assumed that the retailer does not share any demand information with the supplier.

Figure 7 compares the six policies in the (σ_s, σ_I) space for the cases in which the retailer's demand is given by an IID process (left panel) or an MA(1) process (right panel).

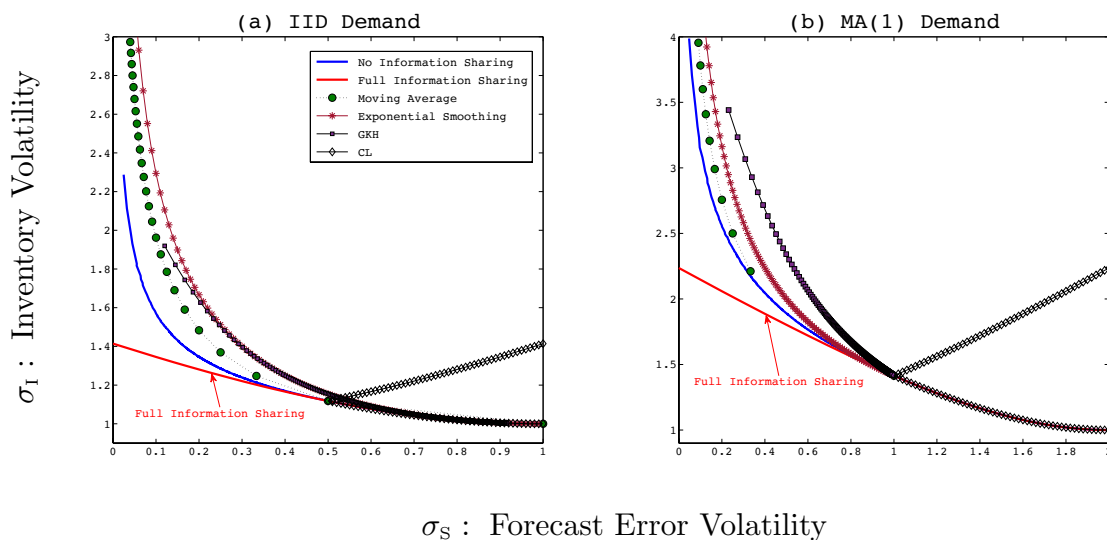


Figure 7 Volatility pairs (σ_s, σ_I) under an optimal replenishment policy (with and without information sharing) and four benchmarks. DATA: In the left panel, $\sigma_\epsilon = \psi_0 = 1$ and $\psi_k = 0$ for all $k \geq 1$. In the right panel, $\sigma_\epsilon = \psi_0 = \psi_1 = 1$ and $\psi_k = 0$ for all $k \geq 1$.

As expected the solution that belongs to the efficient frontier under full information sharing dominates those generated using the (suboptimal) benchmark methods. However, the magnitude of the improvement depends significantly on whether the retailer's order under Full Information Sharing are invertible or not. In the left panel (IID demand) Full Information Sharing orders are invertible when $\sigma_s \geq 0.5$ while in the right panel (MA(1) demand) Full Information Sharing orders are invertible when $\sigma_s \geq 1$. In these regions, not only orders under No Information Sharing but also those generated by the four benchmarks perform very well. On the flip side, when the retailer's orders under No Information Sharing are not invertible (*i.e.*, $\sigma_s < 0.5$ on the left panel or $\sigma_s < 1$

on the right panel) then the performance of the No Information Sharing policy as well as the four benchmarks deteriorate significantly (specially as $\sigma_s \downarrow 0$). This example suggests that (i) the lack of information sharing or (ii) the use of a suboptimal replenishment method can have a significantly negative effect on performance, especially if the retailer wants to minimize the volatility of her orders (i.e., when an optimal ordering policy under full information sharing is not invertible).

Remark 6 *It is worth noticing that the poor performance of some of the alternative smoothing policies in the none-invertibility region is not related exclusively to the lack of information sharing. Indeed, let us consider for instance the exponential smoothing policy. Under full information sharing, the mean forecast error and inventory volatilities are equal to*

$$\sigma_{S|FI}^{ES} = \sigma_\epsilon (1 - \rho) |\psi_0| \quad \text{and} \quad \sigma_I^{ES} = \sigma_\epsilon \sqrt{\sum_{k=0}^{\infty} \left(\sum_{n=0}^k \rho^{k-n} \psi_n \right)^2}.$$

Thus, inventory volatility grows out of bound ($\sigma_I^{ES} \rightarrow \infty$) as mean forecast error goes down to zero ($\sigma_{S|FI}^{ES} \downarrow 0$).

Bullwhip Effect: Next, we look at the magnitude of the *bullwhip effect* under the different policies. Alternative definitions for the bullwhip effect have been used in the literature depending on the interpretation that is used to measure order variability (Chen and Lee 2012) as well as on the data that is available to measure it (Bray and Mendelson 2012). The most popular approach computes the bullwhip effect as the ratio of the order volatility to the demand volatility, that is, under this measure the bullwhip effect in period t is equal to $\mathcal{BW}(t) := \text{Var}(O(t+1))/\text{Var}(D(t+1))$. However, as argued in Chen and Lee (2012), a possibly more appropriate measure –one that better captures the relationship between inventory costs and how orders’ variability propagates through the system– is given by $\widetilde{\mathcal{BW}}(t) := \text{Var}(O(t+1)|\mathcal{M}_s(t))/\text{Var}(D(t+1)|\mathcal{M}_D(t))$. That is, $\widetilde{\mathcal{BW}}(t)$ is an informationally-adjusted measure of the bullwhip effect that emphasizes that the notion of variability that matters is the one that cannot be forecasted. In the context of our model, these two definitions of the bullwhip effect can be written as:

$$\mathcal{BW} = \frac{\sum_{n=0}^{\infty} \tilde{\psi}_n^2}{\sum_{n=0}^{\infty} \psi_n^2} \quad \text{and} \quad \widetilde{\mathcal{BW}} = \frac{\sigma_s^2}{\sigma_\epsilon^2}.$$

We compare these two alternatives measures of the bullwhip effect using the six replenishment policies depicted in panel (b) in Figure 7. To have a uniform (i.e., “apples to apples”) comparison, we compute the bullwhip effect of these policies keeping constant the volatility of the retailer’s inventory. Table 1 depicts the two bullwhip effect measures for the case in which the retailer’s demand is given by an MA(1) process with $\psi_0 = \psi_1 = 1$ and the retailer’s orders are such that the volatility of her inventory is equal to $\sigma_I = 2$.

Table 1 COMPARISON OF BULLWHIP EFFECT MEASURES

	\mathcal{BW}	$\widetilde{\mathcal{BW}}$
Full Information Sharing	1.536	0.072
No Information Sharing	0.812	0.184
Moving Average	1.000	0.232
Exponential Smoothing	0.500	0.250
GKH	0.9215	0.394
CL	1.536	3.000

DATA: Retailer's demand is given by an MA(1) with $\psi_0 = \psi_1 = 1$ and $\sigma_1 = 2$.

A few words about the results in Table 1 are in order. First, note that there is a striking difference between the values of the bullwhip effect computed using these two measures. Furthermore, the different replenishment policies rank quite differently depending on the measure that is used to evaluate the bullwhip effect. For instance, in this example, under the conventional measure \mathcal{BW} the Exponential Smoothing policy minimizes the bullwhip effect while the Full Information Sharing and CL policies have the largest bullwhip effect. On the flip side, if the informationally-adjusted bullwhip measure $\widetilde{\mathcal{BW}}$ is used, then the Full Information Sharing policy would be the one that minimizes the bullwhip effect. It is worth highlighting that the fact that the Full Information Sharing policy minimizes $\widetilde{\mathcal{BW}}$ is not a feature of this particular example but rather a general result since the Full Information Sharing policy is, by construction, the ordering policy that minimizes the supplier's mean squared forecast error for a given value of the volatility of the retailer's inventory.

To the extent that the bullwhip effect is a measure that connects the propagation of orders' variability and system inefficiencies, Table 1 together with the results in the previous sections show that the informationally-adjusted measure $\widetilde{\mathcal{BW}}$ provides a more accurate representation of the bullwhip effect.

6. Illustrative Example: A Two-Tier Supply Chain

Let us now apply the methodology developed in the previous sections in the context of the two-tier supply chain with a single retailer and a single supplier depicted in Figure 1. Our main goal in this section is two-fold: (i) to explicitly connect the volatility measures (σ_s, σ_1) to the retailer's and supplier's operating costs by proposing a concrete model with endogenous cost parameters and (ii) to assess the cost implications of information sharing from the perspective of the retailer, the supplier and the supply chain as a whole.

The specific setting that we consider is a canonical two-tier supply chain system that has been used by various authors (Gavirneni et al. 1999, Lee et al. 2000, or Chen and Lee 2009, among many others) to study similar issues regarding inventory management and supply chain costs under information sharing. On the retailer side, the inventory management problem coincides with the single-stage model discussed in Section 3. That is, we consider a retailer who serves an exogenous demand from a stock of on-hand inventory that she replenishes periodically by placing orders to a single supplier. We assume that the retailer's demand D_t and order O_t in period t satisfy the conditions in Assumptions 1 and 2, respectively. We further assume that demand is fully backlogged and that the supplier offers a guaranteed one-period replenishment lead-time, that is, an order placed by the retailer in period t is received in full in period $t + 1$. It follows that the retailer's on-hand inventory I_t^R at the end of period t evolves according to the dynamics in equation (2). (The superscript 'R' is mnemonic for 'Retailer'.)

In a similar fashion, the supplier serves the retailer's orders using a stock of on-hand inventory that he replenishes by placing production requests every period. We assume the supplier has access to two different modes of production: *normal* and *expedite*. A *normal* production order takes one period to be completed. The supplier is also able to place *expedite* production requests which are delivered in the same period. This expedite mode of production allows the supplier to offer a guaranteed order fulfillment to the retailer. The supplier's inventory I_t^S at the end of period t evolves according to the dynamics: (the superscript 'S' is mnemonic for 'Supplier')

$$I_t^S = I_{t-1}^S + N_{t-1} + E_t - O_t, \quad t = 1, 2, \dots, \quad (23)$$

where N_{t-1} and E_t are *normal* and *expedite* production orders placed by the supplier in period $t - 1$ and t , respectively. The supplier chooses E_t after observing the retailer's order O_t . To avoid trivial solutions, we assume that expedite orders are more expensive to produce than normal ones[†]. It follows that under an optimal production policy $E_t = (I_{t-1}^S + N_{t-1} - O_t)^+$, that is, the supplier only expedites the minimum possible amount to ensure that the retailer's orders are fully satisfied every period.

The supplier chooses his inventory replenishment policy so as to minimize his long-term average expected inventory management cost, which includes regular and expedite production costs as well as inventory holding costs. (Under the full-service requirement and the option of expediting, the supplier never stocks out and does not face any backorder or lost sales costs.) We let C_t^S denote the

[†] Otherwise the supplier's optimal replenishment strategy is trivial, namely, to use only expedite orders for an amount exactly equal to the retailer's orders $E_t = O_t$.

supplier cost in period t , which we assume takes the linear form: $C_t^S := w^N N_t + w^E E_t + h^S I_t^S$, where w^N and w^E are the per-unit normal and expedite production costs (with $w^E > w^N$), respectively, and h^S is the per-unit per-period supplier's holding cost. The supplier's long-term average cost minimization problem is equal to

$$\Pi^S = \inf_{N_t, E_t} \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T C_t^S \right] = \inf_{N_t, E_t} \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T w^N N_t + w^E E_t + h^S I_t^S \right]. \quad (24)$$

Similarly, the retailer's objective is to minimize her long-term average inventory costs that include procurement, holding and backordering costs. We let $C_t^R := w^R O_t + h^R (I_t^R)^+ + b^R (I_t^R)^-$ be the retailer's cost in period t , where w^R is the per-unit procurement cost (or wholesale price charged by the supplier), and h^R and b^R are the retailer's per-unit per-period holding and backordering costs, respectively. The retailer's inventory management problem is as follows:

$$\Pi^R := \inf_{O_t} \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T C_t^R \right] = \inf_{O_t} \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T w^R O_t + h^R (I_t^R)^+ + b^R (I_t^R)^- \right]. \quad (25)$$

Traditionally, the literature on inventory theory and supply chain management has assumed that the cost parameters w^R , h^R , h^S and so on, are exogenously given independent of the firms' inventory management policies. This is certainly a reasonable assumption in the short-term, however, as the planning horizon gets large, one should expect to see a connection between the replenishment policies and these cost parameters. In what follows, we will deviate from this classical modeling paradigm with fixed inventory cost parameters and propose a variation that internalizes the long-term effects of the firms' decisions on their underlying cost structures. In particular, we endogenize the holding cost rates h^S and h^R as well as the retailer's procurement cost w^R by making them dependent on the firms' inventory management policies and in particular the volatility measures (σ_S, σ_I) .

6.1. Endogenous Cost Parameters

Supplier's holding cost h^S : Typically, the financial cost of carrying inventory over time is divided into two main components[‡]: (i) a *physical holding* cost and (ii) an *opportunity* cost. The physical cost refers to those operating expenses associated with handling, warehousing and storing inventory over time. It can also include additional expenses such as security or insurance as well as spoilage and obsolescence if products seat on inventory for a long time. Opportunity costs, on the other hand, are related to the financial returns that are forgone on the funds needed to build and hold

[‡] For example, see Section 6.3 in *Managing Business Process Flows* by Anunpindi et al. (2012).

the inventory for a given period of time. It is a standard accounting practice in inventory theory to express these physical and opportunity costs in a per-unit basis as a fraction of the procurement or production cost. That is, if we denote by w^s the supplier's average per-unit production cost then the cost of holding one unit in inventory during one period of time can be written as $h^s = \alpha^s w^s$, where α^s is the combined per-unit physical holding and capital (or rate of return) cost per period. We will assume that the cost rate α^s is exogenously given reflecting the specific operational and financial conditions under which the supplier operates.

Under the proposed holding cost model, $h^s = \alpha^s w^s$, it becomes clear that different inventory management strategies should lead to different values for h^s in the long run. For instance, if the supplier decides to consistently place large replenishment orders and hold a large amount of inventory then we should see $w^s \approx w^N$. On the other hand, if he decides to satisfy the retailer's demand using exclusively expedite orders (*i.e.*, holds no inventory) then $w^s = w^E$. In order to calculate w^s we need to determine the proportion of the inventory that the supplier procures from each channel. To this end, we make the following assumption:

Assumption 3 (SUPPLIER'S INVENTORY REPLENISHMENT POLICY) *The supplier uses a myopic order-up-to policy to decide the amount of inventory to replenish. In other words, at the end of every period, the supplier decides myopically how much to order for the next period by minimizing the argument inside the summation in (24), conditional on all available information, without internalizing the effect of this decision on future costs.*

A myopic order-up-to policy is in general suboptimal, nevertheless, it has a long tradition in the inventory management literature (*e.g.*, A.F. Veinott 1965) as it offers a simple and tractable approximation to what is, otherwise, a challenging optimization problem to solve in its full generality. Furthermore, an important part of its appeal is that a myopic policy can be optimal under suitable conditions on the underlying demand process (see Johnson and Thompson 1975 for details). It is not hard to see that if the manager were able to order $N_t = Y_t - I_t^s$ in every period then this choice would indeed be optimal. However, unless the supplier were able to return any excess inventory, we need to impose the additional constraint $N_t \geq 0$. The resulting myopic-order-up to policy takes the form $N_t = (Y_t - I_t^s)^+$. This strategy is in general sub-optimal unless $I_t^s \leq Y_t$ (a.s.) for all $t \geq 1$. (For instance, in our model with Gaussian shocks this condition is not satisfied.)

In the statement of the following result, we let \mathcal{L} denote the Loss function

$$\mathcal{L}(z) := \mathbb{E}[(Z - z)^+] = \phi(z) - z(1 - \Phi(z)),$$

where ϕ and Φ are the pdf and cdf of the standard Normal random variable, respectively.

Proposition 4 *Suppose the supplier uses a myopic order-up to policy to replenish inventory. Then, the average per-unit procurement cost w^S is equal to*

$$w^S = w^N + (w^E - w^N) \frac{\sigma_S}{d} \mathcal{L}(z^S), \quad (26)$$

where σ_S is the supplier's mean forecast error defined in equation (9) and z^S is the 'critical fractile'

$$z^S := \Phi^{-1} \left(\frac{w^E - w^N}{w^E - w^N + h^S} \right). \quad (27)$$

Furthermore, the supplier's long-term average cost is equal to

$$\Pi^S = w^N d + h^S \sigma_S z^S + (w^E - w^N + h^S) \sigma_S \mathcal{L}(z^S). \quad (28)$$

At this point, the attentive reader should have noticed that he have used a circular argument in our previous derivation. Indeed, on one hand, the holding cost $h^S = \alpha^S w^S$ is a function of z^S according to equation (26). On the other hand, the value of z^S in the myopic order-up-to policy is a function of h^S according to equation (27). This circularity reflects the fact that in our model the holding cost rate is endogenously determined in an equilibrium that internalizes the interdependence between the inventory strategies and cost structures. Mathematically, h^S solves the following fixed-point condition:

$$h^S = \alpha^S \left[w^N + (w^E - w^N) \frac{\sigma_S}{d} \mathcal{L} \circ \Phi^{-1} \left(\frac{w^E - w^N}{w^E - w^N + h^S} \right) \right]. \quad (29)$$

It is interesting to note that the effect of the retailer's ordering process on the value of h^S is only captured through the square root of the mean squared forecast error σ_S .

Supplier's wholesale price w^R : Let us turn to the connection between the per-unit wholesale price w^R that the supplier charges the retailer and the firms' inventory management policies. It is reasonable to assume that a profit maximizing supplier selects this wholesale price w^R so as to at least covers average per-unit operating expenses. To be concrete, let us assume that the supplier operates under a simple 'cost plus pricing' model under which the supplier charges a wholesale price that is equal to the average per-unit operating cost plus a fixed markup percentage. In other words, the supplier selects w^R so as to achieve a fixed operating margin. Mathematically, let \bar{C}^S be the supplier's average per unit cost then $w^R = (1 + r^S) \bar{C}^S$, where $r^S > 0$ is the supplier's target operating margin. The average cost satisfies $\bar{C}^S = \Pi^S / d$, where Π^S is given in equation (28). As a result, under a *cost plus pricing* model, the wholesale price w^R is equal to

$$w^R = (1 + r^S) \bar{C}^S = (1 + r^S) \left(w^N + h^S \frac{\sigma_S}{d} z^S + (w^E - w^N + h^S) \frac{\sigma_S}{d} \mathcal{L}(z^S) \right). \quad (30)$$

Retailer's holding cost h^R : Using the same line of arguments as the one we use to derive the supplier's holding cost h^S , let us assume that the retailer's long-run average per unit holding cost h^R is proportional to the per-unit procurement cost w^R . That is, $h^R = \alpha^R w^R$ for some fixed positive constant α^R that captures the retailer's physical cost as well as financial (opportunity) cost of holding inventory over time. It follows that

$$h^R = \alpha^R (1 + r^S) \left(w^N + h^S \frac{\sigma_S}{d} z^S + (w^E - w^N + h^S) \frac{\sigma_S}{d} L(z^S) \right). \quad (31)$$

6.2. Retailer's Cost

Equipped with equations (29), (30) and (31) that provide a concrete connection between the retailer's inventory policy and her cost parameters, let us turn to the question of determining an optimal ordering policy. To this end, recall from the analysis in Section 2.1 that the retailer's inventory I_t^R converges in distribution to a Normal random variable I_∞^R with mean \tilde{I}_0 and standard deviation σ_I given in (8). It follows, after some straightforward manipulations, that we can rewrite the retailer's long-term average expected cost as follows:

$$\Pi^R = \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T w^R O_t + h^R (I_t^R)^+ + b^R (I_t^R)^- \middle| \mathcal{M}_D(0) \right] = w^R d + h^R \sigma_I \mathcal{L} \left(-\frac{\tilde{I}_0}{\sigma_I} \right) + b^R \sigma_I \mathcal{L} \left(\frac{\tilde{I}_0}{\sigma_I} \right),$$

where the second equality follows from the weak convergence of I_t^R to I_∞^R and the second equality follows from the definition of the Loss function \mathcal{L} above.

From Remark 3, \tilde{I}_0 is an additional degree of freedom that the retailer can choose to minimize her cost. It follows that an optimal choice is

$$\tilde{I}_0 = \sigma_I \Phi^{-1} \left(\frac{b^R}{h^R + b^R} \right).$$

Combining this condition and the identity $\mathcal{L}(-z) = z + \mathcal{L}(z)$, we conclude that the retailer's long-term average cost is equal to

$$\Pi^R = w^R(\sigma_S) d + \sigma_I \left((h^R(\sigma_S) + b^R) \mathcal{L} \circ \Phi^{-1} \left(\frac{b^R}{h^R(\sigma_S) + b^R} \right) + h \Phi^{-1} \left(\frac{b^R}{h^R(\sigma_S) + b^R} \right) \right).$$

Since both w^R and h^R are increasing functions of σ_S , the retailer can reduce some of her cost components by reducing the value of σ_S and placing orders that are more predictable. This will help the supplier better optimize his own production/inventory costs and pass some of these savings back to the retailer. On the flip side, the retailer's cost also depends on the actual evolution of her

inventory position I_t^R and by minimizing σ_s the retailer can end up increasing this inventory cost component. These are the direct and indirect cost effects that we discuss in the Introduction.

Combining the expressions for σ_I and σ_s derived in Section 3 in terms of the coefficients $\Psi = \{\psi_t\}$ and $\tilde{\Psi} = \{\tilde{\psi}_t\}$ in the MA(∞) representations of D_t and O_t , respectively, together with the endogenous inventory cost model discussed above, the retailer's optimal inventory replenishment strategy solves:

$$\inf_{\tilde{\Psi}} \left\{ w^R(\sigma_s) d + \sigma_I \left((h^R(\sigma_s) + b^R) \mathcal{L} \circ \Phi^{-1} \left(\frac{b^R}{h^R(\sigma_s) + b^R} \right) + h^R(\sigma_s) \Phi^{-1} \left(\frac{b^R}{h^R(\sigma_s) + b^R} \right) \right) \right\} \quad (32)$$

$$\text{subject to } \sigma_I^2 = \sigma_\epsilon^2 \sum_{n=0}^{\infty} (g_n - \psi_n)^2 \quad \text{where } g_0 := 0 \quad \text{and} \quad g_n := \sum_{k=0}^{n-1} (\tilde{\psi}_k - \psi_k), \quad n \geq 1, \quad (33)$$

$$\sigma_s^2 = \sigma_\epsilon^2 \exp \left(\frac{1}{2\pi} \int_{-\pi}^{\pi} \log \left(\tilde{\psi}(e^{-i\lambda}) \tilde{\psi}(e^{i\lambda}) \right) d\lambda \right) \quad \text{where } \tilde{\psi}(z) := \sum_{n=0}^{\infty} \tilde{\psi}_n z^n, \quad (34)$$

$$w^R \text{ and } h^R \text{ are solutions to the system of equations (29), (30) and (31).} \quad (35)$$

Using the results in Section 3 and in particular the representation of the efficient frontier in Definition 2, we can simplify the previous problem in the case in which there is full information sharing. For notational convenience, let us define the retailer's critical fractile

$$z^R(\sigma_s) := \frac{b^R}{h^R(\sigma_s) + b^R}.$$

Then, under full information, the optimization problem (32)-(35) reduces to

$$\min \left\{ w^R(\sigma_s) d + \sigma_\epsilon \sqrt{\psi_0^2 + \left(\frac{\sigma_s}{\sigma_\epsilon} - |\psi_0 + \psi_1| \right)^2} \left(\frac{b^R}{z^R(\sigma_s)} \mathcal{L} \circ \Phi^{-1} (z^R(\sigma_s)) + h^R(\sigma_s) \Phi^{-1} (z^R(\sigma_s)) \right) \right\}$$

$$\text{subject to } 0 \leq \sigma_s \leq \sigma_\epsilon |\psi_0 + \psi_1|$$

$$w^R \text{ and } h^R \text{ are solutions to the system of equations (29), (30) and (31).}$$

This is a simple one-dimensional optimization problem that can be solved numerically very efficiently.

In the case in which there is no information sharing, we are back into the technical issues discussed in Section 4. If the replenishment strategy derived from the solution of the previous optimization problem under sharing is invertible then (by Proposition 3) we know that is also optimal for the case without information sharing. On the other hand, if it is not invertible then we can use the approximation for σ_s in equation (19) to replace constraint (34) and solve numerically the resulting optimization problem.

To illustrate our methodology, let us conclude this section with a set of computational experiments to assess the value of information sharing from the perspective of the retailer, the supplier, and the supply chain as a whole. In these experiments, we use the following numerical values for our model parameters: $\sigma_\epsilon = 1$, $d = 2\sigma_\epsilon$, $w^N = 1$, $w^E = (1 + 25\%)w^N$, $b^R = 25\%w^N$, $\alpha^S = \alpha^R = 30\%$ and $r^R = 50\%$. We note that, without loss of generality, our model allows us to normalize certain parameters as units of measure. For example, in terms of the units of inventory, we have chosen to normalize $\sigma_\epsilon = 1$ and express mean demand d in terms of σ_ϵ . Similarly, we have chosen to normalize the per-unit production cost $w^N = 1$ and express both w^E and b^R in terms of w^N . Figure 8 revisits the example depicted earlier in Figure 7 (panel b) under the specific cost parameters described above and identifies the optimal solution, that is, the optimal volatility pair (σ_S, σ_I) that minimizes the retailer's cost under full information sharing. As we can see, the optimal solution lies in the

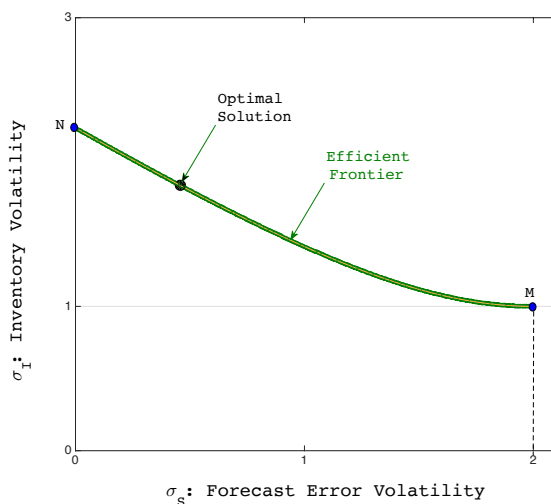


Figure 8 Optimal volatility pair (σ_S, σ_I) under full information sharing. DATA: $\sigma_\epsilon = \psi_0 = \psi_1 = 1$ and $\psi_n = 0$ for all $n \geq 2$.

interior of the efficient frontier meaning that neither the inventory variability nor the forecast error volatility fully dominates the retailer's choice of a replenishment strategy. An interior solution, however, is not always optimal and it is possible that point 'N' or point 'M' become optimal as we vary the cost parameters. For example, let us parametrize the cost of expediting $w^E = (1 + \delta)w^N$ for some $\delta > 0$ (in the previous figure $\delta = 10\%$). Then, we can numerically show that for δ larger than 11.5% the optimal solution switches to point N . On the other hand, the solution moves towards point 'M' asymptotically as $\delta \downarrow 0$.

Let us compare the performance of the supply chain under Full Information sharing (FI) and No Information sharing (NS) to the four benchmark replenishment policies presented in Section 5,

namely, Moving Average (MA), Exponential Smoothing (ES) and the two policies motivated by results in Graves, Kletter, and Hetzel (1998) (GKH) and Chen and Lee (2009) (CL). We also consider as a fifth benchmark the myopic replenishment policy (point 'M' in Figure 8) given its popularity in the inventory management literature.

For each policy type $j = \text{FI, NS, MA, ES, GKH, CL and M}$, we compute the retailer, supplier and supply chain long-term average expected costs Π_j^R , Π_j^S and Π_j^{SC} , respectively. These costs are computed under the assumption that the retailer chooses an ordering policy that minimizes her inventory costs within the specific policy type. For example, the values of Π_{MA}^R , Π_{MA}^S and Π_{MA}^{SC} are calculated under the assumption that the retailer chooses the 'best' Moving Average ordering policy that minimizes her costs. We also assume that with the exception of the FI policy, the retailer is not sharing any demand information with the supplier.

In our computations, we express the performance of this policies relative to Full Information sharing, that is, we define the *relative cost*

$$\text{Relative Cost} = \frac{\Pi_j^k}{\Pi_{\text{FI}}^k},$$

where j ranges over the type of policy $j \in \{\text{NS, MA, ES, GKH, CL and M}\}$ and k ranges over the supply chain agent $k \in \{\text{Retailer, Supplier, Supply Chain}\}$.

Figure 9 depicts the relative cost for the case in which the retailer observes IID demand. Each of the three rows shows a sensitivity analysis on a different parameter of the model. The top row varies the mean demand d , the middle row varies the supplier's cost of expediting w^E and the bottom row varies the retailer's backorder cost b^R . Each column shows the performance of the relative cost for a portion of the supply chain: retailer (left), supplier (center) and entire supply chain (right).

A few remarks about these results are in order. First, in terms of the value of information sharing, we note that lack of collaboration can increase costs quite significant. For instance, on the top row the supplier's costs under a No Sharing policy (NS) are almost 20% higher than those under Full Information sharing when the mean demand d is low[§] while those of the retailer can go up by as much as 10%. The value of information sharing also increases as the supplier's cost of expediting increases (middle row) and seems to be less sensitive to the retailer's backorder cost.

Another point that is worth highlighting is the consistently poor performance of the Myopic policy (M). This is particularly significant given the fact that the Myopic policy is designed to

[§] Under our parameter normalization, a low value of d corresponds to an environment where the retailer demand is highly volatile.

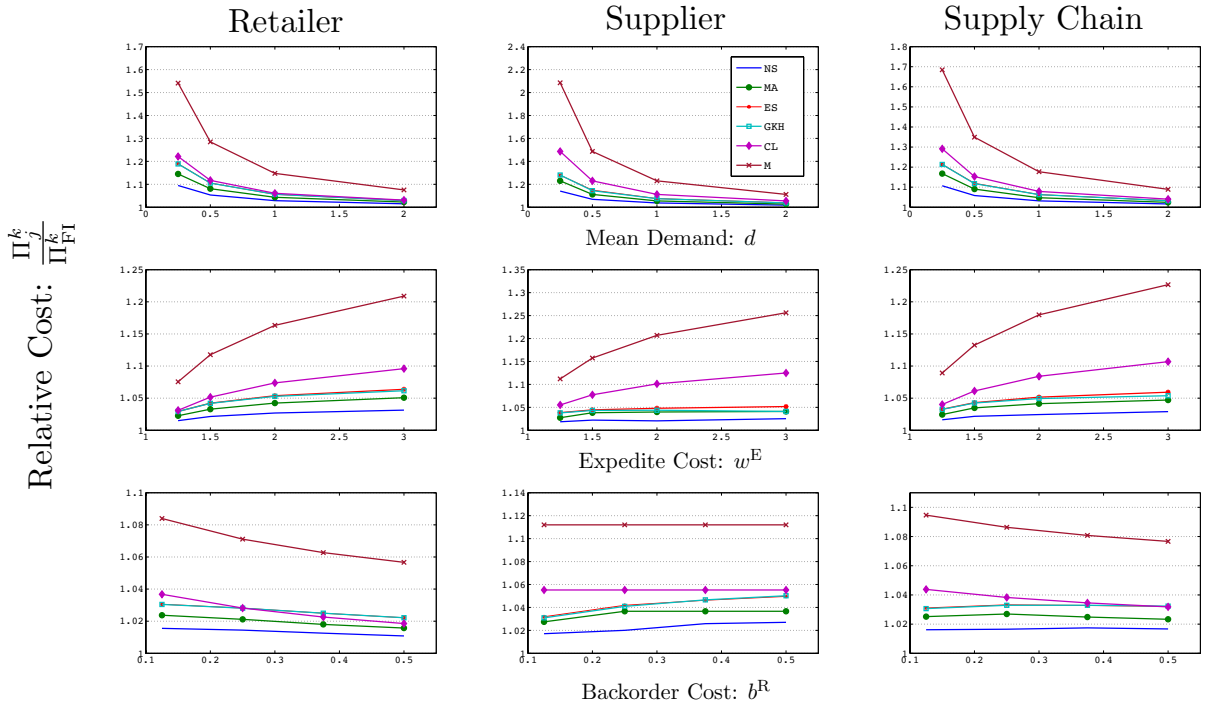


Figure 9 Sensitivity analysis of the *relative cost* of the policies NS, MA, ES, GKH, CL and M for the case in which the retailer observes IID demand. Each of the three rows shows performs a sensitivity on a different parameter of the model: mean demand d (top), supplier's cost of expediting w^E (middle) and retailer's backorder cost b^R . Each column depicts the performance of the relative cost for a segment of the supply chain: retailer (left), supplier (center) and entire supply chain (right).

minimize the variability of the retailer's inventory position, which in an environment where the supplier offers guarantee delivery could be seen as an optimal strategy. However, as we have emphasized before, a myopic policy attempts to minimize the retailer *direct* inventory costs without paying any attention to the *indirect* costs. However, our numerical experiments reveal that these indirect costs can be fairly significant. Actually, in Figure 9 the retailer's costs under a Myopic policy are always 5% or more higher than those under a Full Information sharing policy (and they can be as high as 50% more when mean demand is low).

Finally, in terms of the other benchmark policies (MA, ES, GKH, and CL), it appears that a Moving Average policy offers a slightly better performance than the others.

7. Concluding Remarks

In this paper, we have studied the value of information sharing in the context of a two-tier supply chain system. Our contribution with respect to the existing literature lies in the approach that we

use to connect the value of information sharing to the performance of the system. In particular, we argue that a firm can concentrate the impact of its inventory replenishment strategy into two sources of variability: (a) the variability of its on-hand inventory σ_I and (b) the variability of the forecast error of its orders σ_S . This approach allows us to study in a parsimonious way the structure of an optimal inventory management policy without specifying directly the cost function of the firm.

In the first part of the paper, we characterize the Pareto frontier of efficient pairs (σ_S, σ_I) for the cases in which there is full information sharing and no sharing at all. For the case with full information the Pareto frontier takes a remarkably simple form and depends only on the first two terms of the $MA(\infty)$ representation of the demand. On the other hand, for the case in which there is no information sharing, the analysis is more delicate. We show that if the replenishment order under sharing is invertible then it is also optimal for the case without information sharing. That is, in this case sharing information offers no real value. On the other hand, when the retailer is not able (or willing) to share any demand information we can only provide a partial characterization of an optimal solution and show that information sharing does not always add value. We also show that the question of identifying conditions under which information sharing does offer value reduces to a delicate analysis of the invertibility of a specific stationary process.

In the second part of the paper, we propose a model with endogenous cost parameters that we use to explicitly connect the volatility measures (σ_S, σ_I) with the retailer's operating costs and conduct a set of numerical experiments to illustrate (i) the value of information sharing and (ii) the limitations of using alternative replenishment policies that do not internalize the invertibility or non-invertibility of the retailer's orders.

In terms of future research, our paper leaves open the question of how to compute an optimal replenishment policy under no information sharing when an optimal policy under sharing is not invertible. We conjecture that an alternative parametrization of the problem that works directly with the roots of the characteristic polynomial associated with the retailer's replenishment policy might be better since the issue of checking invertibility becomes trivial in this case. However, more work needs to be done to confirm or refute this claim.

References

- AF Veinott J (1965) Optimal policy for a multi-product, dynamic nonsatiation inventory problem. *Management Science* 12(3):206-222.

- Aviv Y (2001) The effect of collaborative forecasting on supply chain performance. *Management Science* 47(10):1326–1343.
- Aviv Y (2002) Gaining benefits from joint forecasting and replenishment processes: The case of auto-correlated demand. *Manufacturing & Service Operations Management* 4(1):55–74.
- Aviv Y (2003) A time-series framework for supply-chain inventory management. *Operations Research* 51(2):210–227.
- Aviv Y (2007) On the benefits of collaborative forecasting partnerships between retailers and manufacturers. *Management Science* 53:777–794.
- Balakrishnan A, Geunes J, Pangburn MS (2004) Coordinating supply chains by controlling upstream variability propagation. *Manufacturing & Service Operations Management* 6(2):163–183.
- Bray R, Mendelson H (2012) Information transmission and the bullwhip effect: An empirical investigation. *Management Science* 58(5):860–875.
- Brockwell P, Davis R (2006) *Time Series: Theory and Methods* (New York: Springer).
- Chen F (2003) *Supply Chain Management: design, Coordination and Operations*, volume 11 of *Handbooks in OR & MS*, chapter Information Sharing and Supply Chain Coordination (Elsevier B.V.).
- Chen L, Lee H (2009) Information sharing and order variability control under a generalized demand model. *Management Science* 55(5):781–797.
- Chen L, Lee H (2012) Bullwhip effect measurement and its implications. *Operations Research* 60(4):771–784.
- Cui R, Allon G, Bassamboo A, Mieghem JAV (2015) Information sharing in supply chains: An empirical and theoretical valuation. *Management Science* 61(11):2803–2824.
- Gaur V, Giloni A, Seshadri S (2005) Information sharing in a supply chain under arma demand. *Management Science* 51:961–969.
- Gavirneni S, Kapuscinski R, Tayur S (1999) Value of information in capacitated supply chains. *Management Science* 45(1):16–24.
- Giloni A, Hurvich C, Seshadri S (2014) Forecasting and information sharing in supply chains under arma demand. *IIE Transactions* 46:35–54.
- Graves S, Kletter D, Hetzel W (1998) A dynamic model for requirements planning with application to supply chain optimization. *Operations Research* 46(3):35–49.
- Johnson G, Thompson H (1975) Optimality of myopic inventory policies for certain dependent demand processes. *Management Science* 21(11):1303–1307.
- Kovtun V, Giloni A, Hurvich C (2014) Assessing the value of demand sharing in supply chains. *Naval Research Logistics* 61(7):515–531.
- Kurt Salomon Associates (1993) Efficient consumer response: Enhancing consumer value in the grocery industry. Technical report, The Joint Industry Project on Efficient Consumer Response, Washington, D.C.
- Lee H, Padmanabhan V, Whang S (1997) The bullwhip effect in supply chains. *Sloan Management* 38(3):93–102.
- Lee H, Padmanabhan V, Whang S (2004) Information distortion in a supply chain: The bullwhip effect. *Management Science* 50(12):1875–1886.

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

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

Zhang X (2004) Evolution of arma demand in supply chains. *Manufacturing & Service Operations Management* 6:195–198.

Zipkin P (2000) *Foundations of Inventory Management* (Boston: Irwin McGraw-Hill).

Appendix: Proofs

PROOF OF LEMMA 1: Since the ordering policy is admissible (see Definition 1), it follows that $g_L \rightarrow 0$ as $L \rightarrow \infty$. For any integer k , define $f_L = \epsilon_{k-L}$. Using the summation by parts formula, and using the symbol $\lim_{n \rightarrow \infty}$ to denote the limit in mean square as $n \rightarrow \infty$, we have for any fixed integer k

$$\begin{aligned} \sum_{L=0}^{\infty} (\tilde{\psi}_L - \psi_L) \epsilon_{k-L} &= \lim_{n \rightarrow \infty} \sum_{L=0}^n (\tilde{\psi}_L - \psi_L) \epsilon_{k-L} = \lim_{n \rightarrow \infty} \sum_{L=0}^n f_L (g_{L+1} - g_L) \\ &= \lim_{n \rightarrow \infty} [f_{n+1} g_{n+1} - f_0 g_0 - \sum_{L=0}^n g_{L+1} (f_{L+1} - f_L)] \\ &= - \lim_{n \rightarrow \infty} \sum_{L=0}^n g_{L+1} (\epsilon_{k-L-1} - \epsilon_{k-L}) = \sum_{L=0}^{\infty} g_{L+1} (\epsilon_{k-L} - \epsilon_{k-L-1}) . \end{aligned}$$

Therefore,

$$\begin{aligned} \sum_{k=0}^{t-1} r_k &= \sum_{k=0}^{t-1} \sum_{L=0}^{\infty} g_{L+1} (\epsilon_{k-L} - \epsilon_{k-L-1}) = \sum_{L=0}^{\infty} \sum_{k=0}^{t-1} g_{L+1} (\epsilon_{k-L} - \epsilon_{k-L-1}) \\ &= \sum_{L=0}^{\infty} g_{L+1} (-\epsilon_{-L-1} + \epsilon_{t-1-L}) = - \sum_{L=0}^{\infty} g_{L+1} \epsilon_{-L-1} + \sum_{L=0}^{\infty} g_{L+1} \epsilon_{t-(L+1)} \\ &= - \sum_{L=1}^{\infty} g_L \epsilon_{-L} + \sum_{L=1}^{\infty} g_L \epsilon_{t-L} . \quad \square \end{aligned} \tag{36}$$

PROOF OF LEMMA 2: To prove the lower bound we use the fact that $\mathcal{M}_o(t) \subseteq \mathcal{M}_\epsilon(t)$ and so

$$\sigma_s^2 = \text{Var}(O_t | \mathcal{M}_o(t)) \geq \text{Var}(O_t | \mathcal{M}_\epsilon(t)) = \sigma_\epsilon^2 \tilde{\psi}_0^2.$$

On the other side, we can use the identity

$$\log(x) \leq \log(z) + \frac{x-z}{z} \quad \text{for all } x, z \in \mathbb{R}_+$$

in equation (11) to get

$$\begin{aligned} \log\left(\frac{\sigma_S^2}{\sigma_\epsilon^2}\right) &\leq \frac{1}{2\pi} \int_{-\pi}^{\pi} \left[\log(z) + \frac{1}{z} \left(\sum_{n=0}^{\infty} \tilde{\psi}_n^2 + 2 \sum_{n=0}^{\infty} \sum_{k=n+1}^{\infty} \tilde{\psi}_n \tilde{\psi}_k \cos((k-n)\lambda) - z \right) \right] d\lambda \\ &= \log(z) + \frac{1}{z} \left(\sum_{n=0}^{\infty} \tilde{\psi}_n^2 - z \right). \end{aligned}$$

Since the inequality holds for any $z \in \mathbb{R}_+$, we can minimize the right-hand size over z . The minimum is achieved at $z^* = \sum_{n=0}^{\infty} \tilde{\psi}_n^2$. We conclude that

$$\sigma_S^2 \leq \sigma_\epsilon^2 \sum_{n=0}^{\infty} \tilde{\psi}_n^2. \quad \square$$

PROOF OF PROPOSITION 2: Let us consider the optimization problem

$$\begin{aligned} \Sigma_I(\sigma) &:= \inf_{\tilde{\Psi}} \sigma_\epsilon^2 \sqrt{\psi_0^2 + \sum_{n=1}^{\infty} (g_n - \psi_n)^2} \\ \text{subject to} \quad \sigma_\epsilon |\tilde{\psi}_0| &= \sigma, \quad \sum_{n=0}^{\infty} \tilde{\psi}_n = \sum_{n=0}^{\infty} \psi_n, \quad g_0 = 0 \quad \text{and} \quad g_k = \sum_{n=0}^{k-1} (\tilde{\psi}_n - \psi_n), \quad k = 1, 2, \dots \end{aligned}$$

Note first that the first constraint implies that $\tilde{\psi}_0 = \pm\sigma/\sigma_\epsilon$. Let us fix one of these two values for $\tilde{\psi}_0$. Then, the optimization reduces to

$$\begin{aligned} \min_{\tilde{\psi}_1, \tilde{\psi}_2, \dots} &\sum_{n=2}^{\infty} (g_n - \psi_n)^2 \\ \text{subject to} \quad \sum_{n=1}^{\infty} \tilde{\psi}_n &= \sum_{n=0}^{\infty} \psi_n - \tilde{\psi}_0, \quad g_0 = 0 \quad \text{and} \quad g_k = \sum_{n=0}^{k-1} (\tilde{\psi}_n - \psi_n), \quad k = 1, 2, \dots \end{aligned}$$

Note that if we fix $\tilde{\psi}_0$, we also fix the value of g_1 , and that is why the summation in the objective function above starts at $n = 2$. Now, it should be clear that if we can set the values of $\tilde{\psi}_k$ for $k \geq 1$ so that $g_n - \psi_n = 0$ for all $n \geq 2$ then this choice would be optimal. It is not hard to see that this is possible by setting

$$\tilde{\psi}_1 = \psi_2 + \psi_1 + \psi_0 - \tilde{\psi}_0 \quad \text{and} \quad \tilde{\psi}_n = \psi_{n+1} \quad \text{for all } n \geq 2. \quad (37)$$

As a result, the variance of the inventory process is equal to $\sigma_I^2 = \sigma_\epsilon^2 (\psi_0^2 + (\tilde{\psi}_0 - \psi_0 - \psi_1)^2)$. It only remains to decide which of the two feasible values of $\tilde{\psi}_0$ minimizes σ_I^2 . One can see that the optimal choice is $\tilde{\psi}_0 = \text{sign}(\psi_0 + \psi_1) \sigma/\sigma_\epsilon$ and we conclude that

$$\sigma_I = \sigma_\epsilon \sqrt{\psi_0^2 + \left(\frac{\sigma}{\sigma_\epsilon} - |\psi_0 + \psi_1| \right)^2}. \quad \square$$

PROOF OF PROPOSITION 4: Recall that w^N and w^E are the supplier's per-unit production costs for regular and expedited orders, respectively. It follows that the long-term average per-unit production cost is equal to

$$\begin{aligned} w^S &= \lim_{T \rightarrow \infty} \mathbb{E} \left[\frac{\sum_{t=1}^T w^N N_t + w^E E_t}{\sum_{t=1}^T N_t + E_t} \right] = w^N + (w^E - w^N) \lim_{T \rightarrow \infty} \mathbb{E} \left[\frac{\sum_{t=1}^T E_t}{\sum_{t=1}^T N_t + E_t} \right] \\ &= w^N + (w^E - w^N) \lim_{T \rightarrow \infty} \mathbb{E} \left[\frac{\sum_{t=1}^T E_t/T}{\sum_{t=1}^T (N_t + E_t)/T} \right] = w^N + \frac{(w^E - w^N)}{d} \lim_{T \rightarrow \infty} \sum_{t=1}^T \frac{\mathbb{E}[E_t]}{T} \end{aligned} \quad (38)$$

To compute the value of the long-term average expected number of units expedite every period, we need to derive first the myopic order-up-to policy. Under the condition $w^E > w^N$, it follows that at optimality the supplier expedites the minimum possible, *i.e.*, $E_t = (I_{t-1}^S + N_{t-1} - O_t)^-$. Hence, after some straightforward manipulations, we can rewrite the supplier's long-term average cost in equation (24) as follows:

$$\begin{aligned} \Pi^S &= \inf_{N_t, E_t} \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T w^N N_t + w^E E_t + h^S I_t^S \right] \\ &= w^N d + \inf_{N_t} \limsup_T \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T h^S \mathbb{E}[(I_{t-1}^S + N_{t-1} - O_t)^+ | \mathcal{M}_t^S] + (w^E - w^N) \mathbb{E}[(I_{t-1}^S + N_{t-1} - O_t)^- | \mathcal{M}_t^S] \right] \end{aligned} \quad (39)$$

(Recall from Section 2.2 that \mathcal{M}_t^S represents the supplier's linear past in period t .)

Under a myopic order-up-to policy, the supplier decides how much to order for the next period by minimizing the argument inside the summation, that is, by solving every period

$$\min_{N_t} \mathbb{E} [h^S (I_t^S + N_t - O_{t+1})^+ + (w^E - w^N) (I_t^S + N_t - O_{t+1})^- | \mathcal{M}_t^S].$$

Using a change of variable, the optimization problem can be rewritten in terms of an order-up-to level $Y_t = I_t^S + N_t$ that solves

$$\min_{Y_t} \mathbb{E} [h^S (Y_t - O_{t+1})^+ + (w^E - w^N) (Y_t - O_{t+1})^- | \mathcal{M}_t^S].$$

Under Assumptions 1 and 2, O_{t+1} has a normal distribution with mean and standard deviation given by

$$m^S(t) = \mathbb{E}[O_{t+1} | \mathcal{M}_t^S] \quad \text{and} \quad \sigma_S = \sqrt{\text{Var}[(O_{t+1} - m^S(t)) | \mathcal{M}_t^S]} \quad t = 1, 2, \dots$$

(Note that O_t is a stationary process and the standard deviation σ_S is independent of t .) It follows that the optimal order-up to level is equal to

$$Y_t = m_t^S + \sigma_S z_S \quad \text{where} \quad z_S = \Phi^{-1} \left(\frac{w^E - w^N}{w^E - w^N + h^S} \right).$$

The number of units expedited in period $t + 1$ is equal to

$$E_{t+1} = (I_t^S + N_t - O_{t+1})^- = (Y_t - O_{t+1})^- = (O_{t+1} - m_t^S - \sigma_S z_S)^+.$$

Conditional on $\mathcal{M}_S(t)$, the random variable $(O_{t+1} - m_t^S - \sigma_S z_S)$ is normally distributed with mean $-\sigma_S z_S$ and variance σ_S^2 . We can then evaluate $\mathbb{E}[E_{t+1} | \mathcal{M}_S(t)]$ in terms of the Loss function \mathcal{L} to get $\mathbb{E}[E_{t+1} | \mathcal{M}_S(t)] = \sigma_S \mathcal{L}(z_S)$. Plugging this back in equation (38), we get

$$w^S = w^N + (w^E - w^N) \frac{\sigma_S}{d} \mathcal{L}(z_S).$$

Similarly, we can plug the value of Y_t in (39) to get

$$\Pi^S = w^N d + h^S \sigma_S z^S + (w^E - w^N + h^S) \sigma_S \mathcal{L}(z^S). \quad \square$$