Sharing Information to Manage Perishables

Michael E. Ketzenberg*
College of Business
Colorado State University
218 Rockwell Hall
Fort Collins, CO 80523
Tel: (970) 491–7154
Fax: (970) 491–3522
Email:Michael.Ketzenberg@mail.biz.colostate.edu

Mark Ferguson
The College of Management
Georgia Tech
800 West Peachtree Street, NW
Atlanta, GA 30332-0520
Tel: (404) 894-4330
Fax: (404) 894-6030
Email:Mark.Ferguson@mgt.gatech.edu

August, 2005

* Corresponding Author

Sharing Information to Manage Perishables

Abstract

We address the value of information (VOI) sharing in the context of a two–echelon, serial supply chain with one retailer and one supplier that provides a single perishable product to consumers. We evaluate information sharing under two supply chain structures where the retailer shares it inventory level and replenishment policy with the supplier. In the first structure, referred to as Decentralized Information Sharing, both facilities make their own profit-maximizing replenishment decisions. In the second structure, referred to as Centralized Control, the replenishment decisions are coordinated. The latter supply chain structure corresponds to the industry practices of company owned stores or vendor–managed inventory. We measure the VOI as the marginal improvement in expected profits that a supply chain achieves relative to the case when no information is shared. Key assumptions of our model include stochastic demand, lost sales, and fixed order quantities. We establish the importance of information sharing in the supply chain and identify conditions under which relatively substantial benefits are realized. As opposed to previous work on the VOI, the major benefit of information sharing in our setting is driven by the supplier's ability to provide the retailer with fresher product. By isolating the benefit by firm, we show that sharing information is not always Pareto improving for both supply chain partners in the decentralized setting.

Keywords: value of information, vendor managed inventory, supply chain management, perishable inventory

1. Introduction

We explore the value of information (VOI) for inventory replenishment of a perishable product where information may be shared among facilities in a supply chain and used in the decision making to improve performance. As in this study, previous literature on the VOI develops information based replenishment policies and evaluates conditions where shared information is beneficial. Motivation for such research is predicated on the academic prescription that sharing information mitigates the Bullwhip Effect (Lee et al. 1997), that enabling information technologies are widely available, and practitioner initiatives like Efficient Consumer Response (Kurt Salmon Associates 1993). As opposed to previous work on the VOI, the major benefit in our setting comes from the supplier's ability to provide a fresher product.

We place our research in the context of the grocery industry. The importance of perishable goods is growing in terms of sales, SKUs, and the competitive importance of attracting consumers. For supermarkets, perishables are the driving force behind the industry's profitability and represent a significant opportunity for improvement. Perishables account for more than half of supermarket sales in the U.S. or up to $200 billion a year, but also subject the firms to losses of up to 15 percent due to damage and spoilage. Further, the quality, variety and availability of perishables have become an order winning criteria of consumers, representing the primary reason many consumers choose one supermarket over another (Hennessy 1998).

While our research focus is on groceries, the management of perishable inventories is an important problem confronting many other industries including blood banks, food services, pharmaceuticals, chemicals, and increasingly, biotechnology. Yet the grocery industry is particularly appropriate, given the current practitioner activity and industry initiatives. This industry is characterized by a highly competitive business environment with low profit margins, low barriers to entry, the emergence of super centers and mass marketers, increasing consumer demand for high quality perishables, and stagnant industry growth (Saporito 1995). To compete in this competitive environment, many firms are investing in information enabling technologies for the management of their perishables. At the same time, the benefits from these investments remain unclear, as does the distribution of the benefits among the supply chain members.

We address the VOI in the context of a two–echelon, serial supply chain with one retailer and one supplier that provides a single perishable product to consumers. A distinguishing characteristic of perishables is that they have a finite lifetime and hence, the age of the products must be considered in their management. We assume that the product lifetime is fixed and deterministic once produced. Any unsold inventory remaining after the lifetime elapses must be discarded (outdated) at zero salvage value. We evaluate two scenarios. The first scenario, named Decentralized Information Sharing (DIS), considers the case where both supply chain members share their inventory levels and replenishment policies with the other but each facility makes its own profit maximizing replenishment decisions. The second scenario, named Centralized Control (CC), considers the case of coordinated decision making. This second case corresponds to the practice of a company store or vendor–managed inventory (VMI). We formulate the respective scenarios as Markov Decision Processes (MDPs) and measure the VOI as the marginal improvement in expected profit that a supply chain achieves relative to the case when no information is shared. Key characteristics of our model include stochastic demand, lost sales and fixed order quantities.

We establish the importance of information sharing in the supply chain and identify the conditions under which substantial benefits are realized. Through a numerical study, we find that by sharing information, total supply chain expected profits increase an average of 4.2% for a decentralized supply chain and 5.6% for a centralized supply chain. We also show that the benefit of sharing information in the absence of coordination is not always Pareto improving for both firms. Through a sensitivity analysis, we provide conditions where supply chains benefit the most from information sharing and centralized control. In three extensions, we model scenarios where the supplier's revenue is freshness dependent; the fixed lot size is a decision variable; and the consumer demand for the product increases with the freshness of the product.

The rest of the paper is organized as follows: §2 reviews the literature, §3 defines the model, §4 presents our numerical study with discussion, §5 provides extensions, and §6 concludes the paper.

2. Literature Review

Our research draws on two separate research streams: perishable inventory theory and the value of information. In this section, we provide a review of prominent research in each stream and position our study at the point of their intersection.

2.1 Perishable Inventory Theory

The principal distinction within the existing literature on perishable inventory is whether the product has a fixed or random lifetime. We review the key literature on fixed lifetimes since it is more closely related to our research and we refer the reader to Raafat (1991) for a comprehensive review of random lifetimes. Nahmias (1982) provides a good, albeit now dated, literature review of fixed lifetime perishable inventory models.

There are three problems addressed by the literature on fixed lifetime perishable inventory theory: determining reasonable and appropriate methods for issuing inventory, replenishing inventory, and in the case of distribution systems, allocating inventory. Since inventory may contain units of different ages, the issuing problem focuses on the order in which units of each age category are withdrawn from inventory to satisfy demand. Early work by Derman and Klein (1958), Lieberman (1958), and Pierskalla and Roach (1972) collectively show the conditions where issuing the oldest items first (FIFO) and youngest items first (LIFO) are optimal. With constant product utility, as is the case in our model, FIFO issuing is optimal.

Research has also been done to derive and evaluate the optimal replenishment policies for items with a fixed lifetime. Nahmias (1975) and Fries (1975) simultaneously, yet independently were the first to derive and evaluate optimal policies for perishable products with lifetimes greater than two periods. The problem is significantly complicated by the fact that the quantity of inventory in each possible age category must be tracked. They formulate their respective problems as cost–minimizing dynamic programs that include both outdating and shortage costs. In both cases, product is assumed to be fresh on receipt (i.e. fixed lifetime remaining). The optimal policy is shown to be non–stationary and dependent on the age distribution of the inventory. In our model, all units do not arrive fresh at the retailer; the remaining lifetime depends on the age of stock at the supplier used to satisfy a retail order.

Progress on the combined problem of multi–echelon inventory and perishable product inventory systems has been limited. We are aware of only a few contributions in this area, the majority are motivated by the management of blood banks and focus almost exclusively on the allocation problem. Yen (1965), Cohen et al. (1981), and Prastacos (1981) are representative examples. In these studies, replenishment occurs randomly, while in our study it is deterministic.

More recently, Goh et al. (1993) consider a two–stage inventory system at a single facility that is also motivated by the management of blood banks. The first stage contains inventory of fresh blood and the second stage contains older, but still usable, blood. The issuing quantity to the second stage is automatically determined by the age of the blood from the first stage where both the supply and withdrawals of blood occur randomly. Demand requests specify whether they must be satisfied with fresh units or if older units are acceptable.

Fujiwara et al. (1997) provide the most recent contribution to the literature and the only one we are aware of that directly addresses perishable food products. They consider a two–stage inventory system at a single facility where the first stage consists of the whole product (e.g. meat carcasses) that is made up of multiple sub–products (e.g. cuts of meat) while the second stage consists solely of the sub–products. Exogenous demand occurs only at stage two, although unsatisfied stage two demand can be met by emergency issuing from stage one inventory at a cost premium. They derive optimal ordering and issuing policies for this scenario.

Our model extends the research on perishable inventory systems by evaluating a serial system under the assumptions of batch ordering and lost sales: two highly significant and relevant aspects to the management of perishables in the grocery industry.

2.2 Value of Information

The literature on the VOI in a multi-echelon supply chain context is nascent and continues to evolve from the broader literature on multi-echelon inventory systems. There are a few papers that explore the VOI in serial supply chains. Bourland et al. (1996) study how sharing inventory data improves the supplier's ordering decisions with stationary stochastic demand. In their model, the VOI manifests itself in the supplier's ability to respond to the change in the retailer's inventory level, prior to the placement of the retailer's order.

Chen (1998) compares echelon stock policies that require information sharing and centralized decision making with installation–stock policies that do not require information sharing and allow independent decision–making. Although he reports a cost improvement with an echelon policy by as much as 9%, on average the benefit is reported at 1.8%.

Gavirneni et al. (1999) explore the impact of a supplier's capacity restriction on the VOI. They develop two cases of information sharing: 1) the retailer shares information about underlying demand and the parameters of its order policy and 2) the retailer also communicates its inventory level. They report a high level of VOI in the first case, but only an incremental additional benefit from sharing its inventory level in the second case.

Lee et al. (2000) address the VOI when demand follows an AR(1) process and is correlated one period to the next. They show that sharing demand information can lead to substantial benefits, particularly when demand correlation is high. Raghunathan (2001), however, points out that the supplier's base stock policy used in Lee et al. (2000) without information sharing only utilizes the last observed order from the retailer. He shows that when the full history of orders is used, the VOI is negligible. Other studies investigate the VOI in the context of distribution systems consisting of one supplier and N retailers. Examples include Cachon and Fisher (2000), Aviv and Federgruen (1998), and Moinzadeh (2002).

Beyond our study, Ferguson and Ketzenberg (2005) is the only study we are aware of that addresses the value of information sharing in the context of perishable inventory. The authors address the value of information sharing from a supplier to one of its many smaller retailers. In their study, the supplier shares its age–dependent inventory state, replenishment policy, and demand information with the retailer. In contrast, we examine the reverse flow of information where the retailer shares information with the supplier. Also, Ferguson and Ketzenberg (2005) model a retailer in a large distribution network where the supplier's ordering policy is not dependent on a single retailer's actions whereas we model a serial supply chain where the single retailer's actions are highly relevant to the supplier's decisions.

3. Model

The setting is a serial supply chain consisting of two echelons, a retailer and supplier that provide a single perishable product to consumers that has a deterministic lifetime of $M + 1$ periods. Throughout its lifetime, the utility of the product remains constant until the remaining lifetime is zero periods, after which the product expires and is outdated (disposed) without any salvage value. This assumption corresponds to the wide–spread use of product expiration dates on packaged goods such as fresh cut meat, dairy products, and packaged produce.

We assume a periodic review inventory model for each facility, as this is the most common system used in the grocery industry. For the retailer, the order of events each day follows the sequence: 1) receive delivery, 2) outdate inventory, 3) place order, and 4) observe and satisfy demand. Retail demand is discrete, stochastic, and stationary over time. Let D denote total demand in the current period, with probability mass function $\phi(\cdot)$, mean μ, variance σ^2, and C the corresponding coefficient of variation. Unsatisfied demand is lost. To simplify notation, we normalize the retailer's revenue per unit of satisfied demand to one dollar and predicate the unit purchase cost on the product margin m_0, expressed as a percentage of unit revenue. A holding cost h_0 is assessed on ending inventory.

The replenishment decision q is restricted to either zero or Q units, where the batch size Q represents the bundle of units that are packaged, shipped, and sold together. The fixed batch size Q captures certain economies of scale in transportation and handling. The assumption of a fixed batch size is common in the literature on the value of information (e.g., see Moinzadeh

2002, Cachon and Fisher 2000, Chen 1998). A fixed batch size is also commonly observed in practice for many perishable items. Grocery stores commonly replenish when the inventory position is down a case. Also, because of increasing levels of product variety there are thousands of low volume products where a single batch of replenishment is sufficient to satisfy expected demand during the order cycle. Specifically, our restriction on the order size enables us to track the age of product as it moves between echelons, a key modeling contribution to the literature. In a later section we show how our model can also be used to find an optimal value of Q.

The replenishment lead-time is one period. Since the product is perishable, inventory may be composed of units with different ages. Let i_x denote inventory, after outdating and before demand, that expires in x periods, where $x = 1, \ldots, M$ and M is the maximum product shelf life at the retail echelon. Let $\vec{i} = (i_1, i_2, \ldots, i_M)$ represent the vector of inventory held at each age class and define $I = \sum_{x=1}^{M} i_x$. Demand is satisfied using a FIFO inventory issuing policy and inventory is not capacitated.

For the supplier, the order of events each period follows the sequence: 1) receive delivery, 2) observe and satisfy demand, and 3) place order. An order placed by the retailer corresponds to a demand at the supplier in the same period. Since the supplier only observes orders of Q units and faces no ordering cost, the supplier replenishes in orders of Q units. We assume that the supplier orders from a perfectly reliably exogenous source (i.e. the outside source has ample capacity) and the lead-time is one period (i.e. whenever Q units are ordered they become available at the start of the next period). Thus, the supplier faces uncertainty only in the timing of the order arrivals. If the supplier receives an order and does not have units in stock to fulfill it, the supplier pays an expediting charge that allows it to meet the order in the same period. Thus, the retailer always receives its order request one period after placing it.

The supplier's replenishment policy corresponds to a time phased order point policy incorporating safety lead-time. Denoted by α , safety lead time represents the number of periods the supplier waits after receiving a retailer order before it places its own replenishment order so that $\alpha \in (0, 1, ..., M)$. The safety lead-time is based on the supplier's critical fractile, determined from its cost of being early or late with a replenishment order. This policy is optimal for a firm facing intermittent demand with deterministic quantities, uncertain timing, and non–perishable inventory (Silver et al. 1998). It is also the optimal policy for the supplier since no outdating occurs at the supplier's location. This is because the longest possible time between retail orders is M periods and, at that time, the age of product at the supplier has a minimum life of two periods remaining. Note that this assumption requires a further condition that the retailer will never intentionally go through a period with zero inventory, thus assuring the interval between retail orders never exceeds M periods. Although restrictive, our assumptions are supported by industry where 1) outdating at supplier echelons is trivial compared to the retail echelon and 2) there exists a strong emphasis on high retail in-store availability.

3.1 No Information Sharing (NIS) Case

We begin by establishing a base case where the retailer does not periodically share information pertaining to its replenishment process or inventory position. We formulate the retailer's Markov Decision Process (MDP) and the supplier's corresponding replenishment policy assuming the supplier only observes the timing between the retailer's orders.

3.1.1 NIS Case: Retailer's Policy

We formulate the retailer's replenishment problem as a MDP where the objective is to find an optimal reorder policy that maximizes expected profit. The linkage between periods is captured through the one period transfer function of the retailer's age dependent inventory. This transfer is dependent on the current inventory level, any order placed in the current period, the realization of demand *D* in the current period, and the remaining lifetime of any replenishment inventory (this determines the position *x* within the vector $\vec{i}$ that is updated with the replenishment quantity). The remaining lifetime of replenished inventory, denoted as *A,* is a function of the number of periods since the last retailer order L where $A, L \in \{1, 2, ..., M\}$, and the supplier's safety lead–time α (described more fully in Appendix A).

For ease of exposition, let $(z)^+ \equiv \max(z, 0)$ and z' denote a variable defined for the next period, whereas a plain variable *z* is defined for the current period. Let $\vec{i}'$ denote the retailer's inventory level in the next period and $\tau(\vec{i}, D, q, A)$ denote the one period transfer function. Then $\vec{i}' = \tau(\vec{i}, D, q, A)$ where

$$i'_x = \begin{cases} \left(i_{x+1} - \left(D - \sum_{z=1}^{x} i_z \right)^+ \right)^+ & if \ 0 < x < A \\ q & if \ x = A \end{cases}.$$

Now, let $G(I)$ denote the retailer's one period profit function where

$$G(I) = \sum_{D=0}^{\infty} \left[\min(D, I) - h_0 (I - D)^+ \right] \phi(D).$$

We now introduce the retailer's MDP. The value $\bar{c}$ is the equivalent average return per period when an optimal policy is used. The extremal equations are

$$f\left(\vec{i},L\right)+\overline{c}=\max_{q\in\{0,Q\}}\left\{G\left(I\right)-q\left(1-m_0\right)+\sum_{D=0}^{\infty}f\left(\tau\left(\vec{i},D,q,A\right),L'\right)\phi\left(D\right)\right\} \quad (1)$$

where

$$A=\begin{cases}M & \text{if } L\le\alpha\\ M-L+\alpha+1 & \text{if } L>\alpha\end{cases} \quad (2)$$

$$L'=\begin{cases}1 & \text{if } q=Q\\ L+1 & \text{if } q=0\end{cases} \quad (3)$$

Since the state and decision spaces are discrete and finite and profit is bounded, there exists an optimal stationary policy that does not randomize (Putterman, 1994 pg 102 - 111). The left hand side of (1) defines an extremal equation by the vector of inventory $\vec{i}$ and the number of periods L since the last order was placed. The right hand side of (1) computes the total expected profit composed of the one period profit function, the purchase cost associated with any new order, and future expected profit. Equation (2) determines the remaining lifetime of any receipts. Note that if $L\le\alpha$, then $A=M$ since replenishment occurs through expediting. Also, (2) assumes that the retailer knows both the supplier's safety lead–time α and the age of replenishment A. The retailer can readily deduce these values given the replenishment history with the supplier. Finally, (3) updates the number of periods since the last order was placed, predicated on whether or not an order is placed in the current period.

3.1.2 NIS Case: Supplier's Policy

Since the retailer is restricted to ordering Q units at a time, the supplier also replenishes in batch sizes of Q units. A sample path of the supplier's inventory level follows a renewal process with the renewal occurring each time the retailer places an order. The supplier's objective is to make an ordering decision that minimizes its inventory related cost over this renewal cycle.

Since the supplier is only concerned with the timing of its replenishment, the problem reduces to a myopic cost minimization problem that the supplier faces each period it ends with zero units in inventory. If the supplier does not have inventory when the retailer places an order, the supplier pays an expediting charge of *b*. If the supplier does have inventory and the retailer does not order, the supplier pays a holding cost of h_1 for each of the Q units it holds. Obviously, if the number of periods since the retailer's last order is equal to the lifetime of the product then the supplier knows that the retailer will place an order the next period. Let β represent the number of periods since the retailer's last order and λ represent the supplier's decision to place a replenishment order (from his own supplier) where $\lambda = 1$ corresponds to an order being placed and $\lambda = 0$ corresponds to a decision not to order. Also, let q' denote the retailer's expected order in the next period. The supplier's expected inventory related cost for the next period depends on the conditional probabilities that the retailer places an order next period $P(q' = Q \mid \beta)$ or not $P(q' = 0 \mid \beta)$ given that the number of periods since the last retailer order will be β in the next period. If $\beta = M$, the supplier knows with certainty the retailer will place an order next period. The supplier's objective is

$$\min_{\lambda \in \{0,1\}} \; (1-\lambda) b * P(q' = Q \mid \beta) + \lambda (h_1 Q * P(q' = 0 \mid \beta)) .$$

It is straight–forward to show that the supplier's optimal policy is

$$\lambda^* = \begin{cases} 1 & \textit{if} \;\; \beta = M \\ 1 & \textit{if} \;\; P(q' = Q \mid \beta < M) \geq \frac{h_1 Q}{b + h_1 Q} \\ 0 & \textit{otherwise} \end{cases} .$$

In Appendix A, we characterize the probability $P(q' = Q \mid \beta)$. Note that we assume the supplier acts honorably and does not attempt to increase its profit by ordering earlier than the safety lead-time so that the product's useful life at the retailer will be shorter, forcing the retailer to order more frequently. While there may be a short-term incentive for the supplier not to act in this manner, the long-term negative consequences do not typically make it worthwhile, as the retailer would eventually figure out the supplier's deceitfulness.

To express the supplier's expected profit per period, some additional notation is required. Let $\pi_{\vec{i},L}$ denote the steady state probability that the retailer is in state $\left(\vec{i},L\right)$ and let $q^*_{\vec{i},L}$ denote the retailer's corresponding optimal replenishment decision for this state. Further, let m_1 denote the supplier's margin per unit expressed as a percentage of its unit revenue. The supplier's expected profit per period is

$$\sum_{\vec{i}}\sum_{L}\begin{cases}\left[m_1\left(1-m_0\right)q^*_{\vec{i},L}-b\right]\pi_{\vec{i},L} & if\ L-\alpha\leq 0\ and\ q^*_{\vec{i},L}>0\\ \left[m_1\left(1-m_0\right)q^*_{\vec{i},L}-h_1\left(Q-q^*_{\vec{i},L}\right)\right]\pi_{\vec{i},L} & if\ L-\alpha>0\\ 0 & otherwise\end{cases}.$$

3.2 Decentralized Information Sharing (DIS) Case

In the DIS Case the retailer shares its inventory state and replenishment policy with the supplier, but decision-making remains independent. As before, we start by formulating the retailer's MDP and then express the supplier's replenishment policy.

3.2.1 DIS Case: Retailer's Policy

The retailer's optimization is similar to the NIS Case except that it is now necessary to track the supplier's inventory state since the supplier's replenishment decision is now state–dependent on the retailer's inventory position. To reduce notational complexity, we track the supplier's age dependent inventory by using A – the remaining *retail* shelf life, since the age at the supplier is simply $A+1$ if the supplier holds inventory. This involves a slight change in interpretation, since now A takes values in $\{0,1,...,M\}$ and $A=0$ corresponds to the state when the supplier has zero inventory and implicitly the age of replenished items will be M due to expediting. Since we now track the supplier's inventory with A, we drop L (the periods since the last retailer order) from the state space. The extremal equations are

$$f\left(\vec{i},A\right)+\bar{c}=\max_{d\in\{0,Q\}}\left\{G\left(I\right)-q\left(1-m_0\right)+\sum_{D=0}^{\infty}f\left(\tau\left(\vec{i},D,d,A\right),A'\right)\phi\left(D\right)\right\} \tag{4}$$

where

$$A'=\begin{cases}A-1 & if\ \lambda=0 \text{ and } q=0\\ M & if\ \lambda=1\\ 0 & otherwise\end{cases}. \tag{5}$$

Note that (5) determines the supplier's inventory state in the next period, predicated on both the retailer's order and the supplier's replenishment decision. In the next section, we describe the supplier's policy that incorporates the information shared by the retailer.

3.2.2 DIS Case: Supplier's Policy

Under the DIS Case, the supplier's objective is

$$\min_{\lambda\in\{0,1\}} (1-\lambda)(b*P(q'=Q|\vec{i}))+\lambda(h_1Q*P(q'=0|\vec{i}))$$

which gives an optimal policy of

$$\lambda^*=\begin{cases}1 & if\ P(q'=Q|\vec{i})\geq\frac{h_1Q}{b+h_1Q}\\ 0 & otherwise\end{cases}.$$

The conditional probabilities $P(q'=Q|\vec{i})$ and $P(q'=0|\vec{i})$ are functions of the retailer's one-period inventory state transition probabilities and the optimal ordering decisions resulting from (4). Since the retailer and supplier replenishment decisions are inter-related and decision–making is decentralized, some discussion is warranted regarding the order in which the values for q^* and λ^* are determined. We employ the following solution procedure: 1) Given a system state $\left(\vec{i}, A\right)$, condition on the decision $q = 0$ and compute the optimal supplier policy $\lambda^* | q = 0$. 2) Compute the corresponding expected average profit for the retailer given these decisions. 3) Provide the same treatment to the condition for the decision $q = Q$ and find both the optimal supplier policy $\lambda^* | q = Q$ and the associated expected average profit for the retailer. 4) Choose the value q^* that maximizes the retailer's expected profit.

As in the NIS Case, the supplier's expected average profit per period is determined from the limiting behavior of the retailer in steady state. Letting $\pi_{\vec{i},A}$ denote the steady state probability that the system is in state $\left(\vec{i}, A\right)$ and $q^*_{\vec{i},A}$ denote the corresponding optimal retailer replenishment decision, the supplier's expected profit per period is

$$\sum_{\vec{i}} \sum_{A} \begin{cases} \left[m_0 \left(1 - m_1\right) q^*_{\vec{i},A} - b \right] \pi_{\vec{i},A} & \text{if } A = 0 \text{ and } q^*_{\vec{i},A} > 0 \\ \left[m_0 \left(1 - m_1\right) q^*_{\vec{i},A} - h_1 \left(Q - q^*_{\vec{i},A}\right) \right] \pi_{\vec{i},A} & \text{if } A > 0 \\ 0 & \text{otherwise} \end{cases} .$$

3.3 Centralized Control (CC) Case

In the CC Case, a central decision maker seeking to maximize total supply chain profits makes replenishment decisions for both the retailer and the supplier. This corresponds to the practice of vendor–managed inventories (VMI). The retailer no longer places orders with the supplier. Instead, we interpret the decision variable q as a planned shipment from the supplier to the retailer. In addition, the supplier's replenishment order λ is now added to the decision space of the MDP. Theoretically, if it is optimal for the supplier to place an order in a period where it already has Q units in inventory, the existing inventory is immediately disposed. We note, however, that this scenario has never occurred in our numerical studies. For convenience, let $c_1 = Q\left(1 - m_0\right)\left(1 - m_1\right)$ denote the supplier's purchase cost. The extremal equations are

$$f\left(\vec{i},A\right)+\overline{c}=\max_{q\in(0,Q),\lambda\in\{0,1\}}\left(\begin{array}{l} G(I)-c_1\lambda+\sum_{D=0}^{\infty}\left[f\left(\tau\left(\vec{i},D,q,A\right),A'\right)\right]\phi(D) \\ -\begin{cases} 0 & if\ A=0\ and\ q=0 \\ c_1-b & if\ A=0\ and\ q>0 \\ h_1(Q-q) & otherwise \end{cases} \end{array}\right). \qquad (6)$$

Since the objective is to maximize system–wide profit, the optimization expressed in (6) omits the transfer price between the supplier and the retailer. Instead, expected profit maximized in the MDP is the sum of the one period profit function, the purchase cost to the supplier for regular replenishment, the purchase cost plus penalty cost for any supplier expediting, holding costs applied to ending inventory for both facilities, and future expected profit. Note that (5) carries–over from the DIS Case and is not repeated here.

4. Numerical Study

We evaluate the VOI in the DIS Case and the Value of Centralized Control (VCC) in the CC Case where VOI and VCC are measured as the % improvement in expected total supply chain profit, relative to the NIS Case. Specifically,

$$VOI=\frac{\left(E\left[\text{Profit}_{DIS}\right]-E\left[\text{Profit}_{NIS}\right]\right)}{E\left[\text{Profit}_{NIS}\right]} \text{ and } VCC=\frac{\left(E\left[\text{Profit}_{CC}\right]-E\left[\text{Profit}_{NIS}\right]\right)}{E\left[\text{Profit}_{NIS}\right]}.$$

Consumer demand $\phi(\cdot)$ corresponds to a truncated negative binomial distribution with a maximum value of 50 (any probabilities for demand exceeding 50 are redistributed proportionately within the truncated limit of the distribution). See Nahmias and Smith (1994) regarding the advantages of assuming negative binomial distributions for retail demand. Across our numerical experiments, the mean of the distribution is held constant at four and the Coefficient of Variation (*C*) is treated as a parameter to the model using the values reported below. Each period represents a day and the holding cost at each echelon is 40% of the purchase cost, measured on an annual basis. In total, we consider 972 experiments that comprise a factorial design for all combinations of the following parameters:

$$C \in (0.5, 0.6, 0.7) \quad M \in (4, 5, 6) \quad Q \in (8, 9, 10) \quad m_0 \in (0.4, 0.5, 0.6)$$

$$m_1 \in (0.4, 0.5, 0.6) \quad b \in (0.05c_1, 0.10c_1, 0.15c_1, 0.20c_1)$$

Our selection of parameter values is based on values observed in practice for several common products like fresh meat and seafood, deli items, ready-made meals, and packaged produce. At the same time, our selection is constrained by the computational feasibility of the resulting MDP, since the size of the state space expands exponentially with the vector of age–dependent inventory. Current and available computing technology enables us to solve a MDP of approximately twenty million states in twenty minutes. Notwithstanding, the range of parameter values considered covers an extensive selection of products (Office of Technology Assessment Report, 1979).

For each experiment, we use value iteration to compute the results for the respective MDPs and then solve the accompanying state transition matrices using the method of Gaussian elimination to evaluate steady state behavior as described in Kulkarni (p. 124). In §4.1, we discuss our general observations and in §4.2 we report the results of our sensitivity analysis.

4.1 Computational Results and General Observations

In general, we find that both information sharing and centralized control can lead to a considerably fresher product for sale at the retailer. In Table 4.1, we report the VOI for the entire supply chain and for each member under a decentralized structure (DIS Case) and the corresponding VCC for the total supply chain under a centralized structure (CC Case), at given percentiles of the 972 experiments. For example, the 0.50 percentile denotes the median values for VOI and VCC. From this table, three observations emerge: 1) the VOI is lower than the

VCC, although it can still be substantial, 2) the VOI is not necessarily shared equally between the retailer and the supplier, and 3) both the VOI and the VCC are sensitive to model parameterization and depend largely on system behavior as we discuss for each case below.

	DIS Case			CC Case
Percentile	Total	Retailer	Supplier	Total
0.00	0.0%	0.0%	-10.1%	0.0%
0.25	0.8%	1.2%	-1.6%	1.2%
0.50	3.3%	4.1%	0.3%	4.6%
0.75	7.0%	10.1%	4.8%	8.7%
1.00	13.3%	26.9%	19.0%	16.0%
Mean	4.2%	6.2%	1.6%	5.6%

Table 4.1: VOI (DIS Case) and VCC (CC Case) across experiments

4.1.1 DIS Case Observations

In the DIS Case, information sharing enables the supplier to better time the arrival of its replenishment with the timing of retail orders. In turn, the freshness of product (measured in terms of the expected average lifetime remaining) replenished at the retailer increases from an average of 3.77 periods to 4.46 (18.3% increase). Thus, product outdating at the retailer decreases by an average of 39.0%. This increased product freshness also enables the retailer to boost its service level by 3.1% on average.

The change in retailer performance has two direct effects on the supplier that are related to a change in the retailer's average per period order quantity to the supplier. The change reflects both a decrease in outdating at the retailer and an increase in retailer service. When the increase in retailer service (and hence units of satisfied demand) exceed the reduction of outdating, the supplier realizes a net increase in retailer orders and the supplier is better off. The converse results in a net decrease in retailer orders and the supplier is worse off. Across experiments, we find that half of the time the combination results in a net decrease in retailer orders and this can be as large as a 10.5% reduction. In the other half of the experiments, a net increase in retailer orders occurs and this can be as large as an 18.5% increase. It is important to note that even though the supplier is able to reduce its expected inventory related costs in all experiments through information sharing; these savings are generally trivial compared to the increase or decrease in revenue that arises through the change in retailer behavior. In section §4.2 we evaluate the conditions which affect the retailer's order stream in a sensitivity analysis.

Total supply chain profit always improves with information sharing, even when the supplier's profit decreases. An important avenue for future research is to explore how certain contracts and incentives can be implemented so that the maximum benefits from information sharing can be realized and Pareto improving. Without such arrangements, it is doubtful the supplier will be a willing participant.

4.1.2 CC Case Observations

With centralized control, the improvement in total supply chain profit is greater than the improvement observed with information sharing. On average, the VCC is 27% greater than the VOI. There are two effects at work here. First, there is minimal value in holding inventory at the supplier. Thus, the supplier serves a cross–docking function wherein any replenishment it receives is immediately sent onward to the retailer. We observe an average decrease of 44% in the supplier's expected inventory holding costs and a related average improvement of 24% in the freshness of the product delivered to the retailer. Note that this represents over a 5% improvement in product freshness relative to the DIS Case.

The second effect comes from the elimination of double marginalization (the stocking decision at the retailer is predicated on the entire supply chain's profit, not just the retailer's as in the NIS and DIS Cases). Consequently, the retailer's service level increases an average of 7.0%. This represents a considerable improvement when compared to the 3.1% average increase observed in the DIS Case. To provide the higher level of service, a higher level of inventory is positioned at the retailer and therefore the system may experience an increase in outdating relative to both the NIS Case and DIS Case.

4.2 Sensitivity Analysis

In this section, we provide a sensitivity analysis on the VOI and the VCC. Generally, we find that both are sensitive to product perishability, the retailer's ability to match supply and demand, and the size of the penalty for mismatches in supply and demand. We illustrate the sensitivity of the VOI and the VCC to each parameter in Figure 4.1. The height of each bar corresponds to the average VOI/VCC across experiments for the parameter value specified on the x-axis. We discuss these relationships below. For reference, we also provide a more complete set of performance measures in Appendix C.

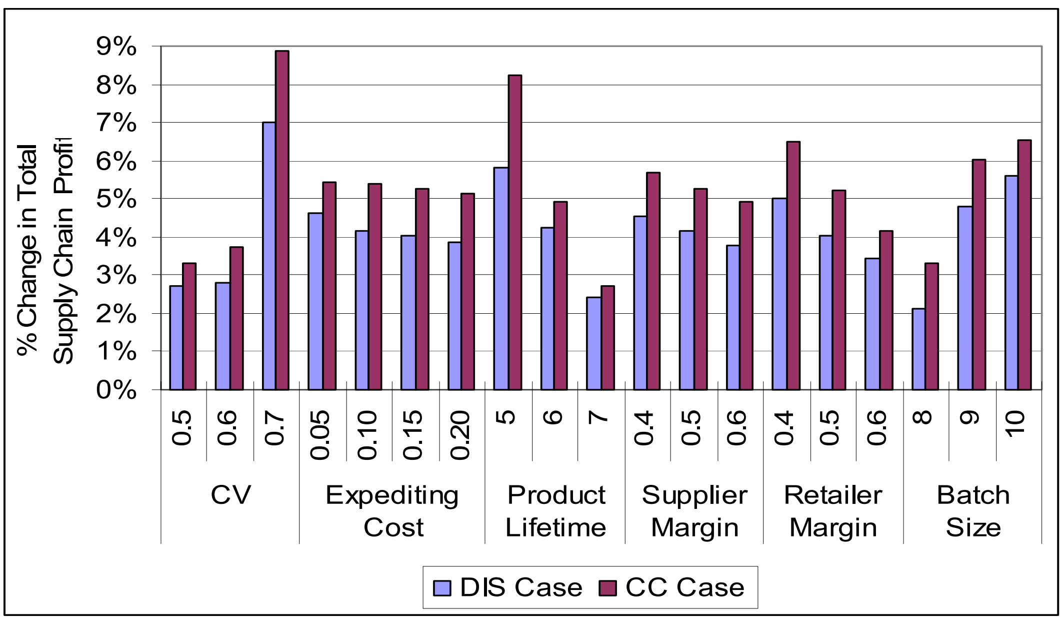

Figure 4.1: Sensitivity of the average VOI/VCC for each fixed parameter value

4.2.1 Product Perishability

As shown in Figure 4.1, the VOI and the VCC decrease with respect to increases in the product lifetime. The main benefit of information sharing is the supply of fresher product to the retailer. When the product lifetime is short, improvements in product freshness have a larger impact on the retailer's service level than when the product lifetime is long. Fresher product reduces the potential for outdating, allowing the retailer to carry more inventory for the same amount (or less) of product outdating which results in higher sales so that the entire supply chain is better off. However, the VOI and the VCC does not always increase with decreases in the product lifetime, as both the product lifetime and batch size impose constraints on the supplier's ability to improve product freshness. Certainly for a product lifetime of one day, the replenishment problem reduces to a newsvendor problem and there is no value with respect to information sharing. In Figure 4.2, we show through an illustrative example that the VOI and the VCC are actually concave with respect to the product lifetime. Here we vary $M \in (2,3,4,5)$ with a fixed set of parameter values: $\mu = 4$, $C = 0.7$, $Q = 7$, $b = 0.2c_1$, $m_1 = 0.5$, and $m_0 = 0.6$.

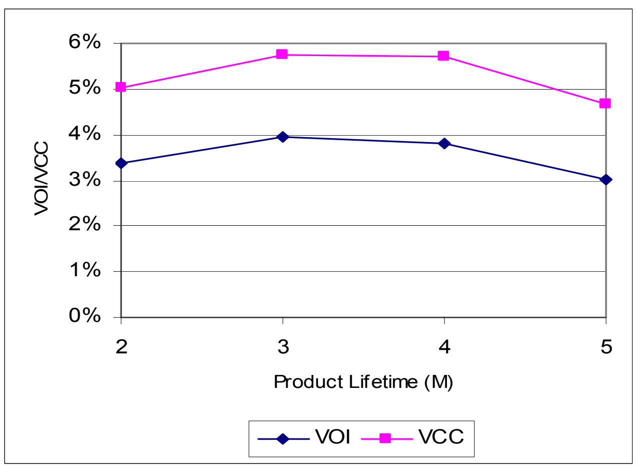

Figure 4.2: Sensitivity of the average VOI/VCC for product lifetime

Long product lifetimes result in small VOI and VCC since the prospect of outdating is small. In this scenario, service levels are higher and outdating is lower so that any improvement in product freshness will not materially change the retailer's behavior. To see this, consider the extreme case of a non-perishable product. Here, there is no outdating and the only benefit of information sharing is to improve the supplier's ability to minimize its own related inventory costs which typically represent a small portion of total supply chain costs. To demonstrate, we duplicate our experimental design (excluding variation with respect to the product lifetime) for the case of non-perishable products. In total, there are 324 experiments and we find in all cases that both the VOI and the VCC are trivial: the average is 0.1% and the maximum is 1.3%.

4.2.2 Matching Supply and Demand

Two factors that affect the retailer's ability to efficiently match supply with demand are demand uncertainty, measured as the coefficient of variation C, and the batch size Q. As shown in Figure 4.1, it is clear that as these parameter values increase, so does the VOI and the VCC. The more difficult it is for the retailer to match supply with demand, the more perishability becomes an issue. We further validated our result with respect to Q by examining the VOI and the VCC for smaller batches sizes, $Q \in (5, 6, 7)$, than those in our main study. We report the results in Table 4.2 where the values for VOI and VCC are averaged across experiments at each level of M and Q. It is quite clear that both the VOI and the VCC quickly approach zero as the batch size approaches the mean demand rate.

		Retail Lifetime		
	VOI	4	5	6
	5	1.2%	0.1%	0.0%
Batch Size	6	1.4%	0.2%	0.1%
	7	1.4%	0.4%	0.2%

		Retail Lifetime		
	VCC	4	5	6
	5	1.5%	0.1%	0.0%
Batch Size	6	2.3%	0.6%	0.0%
	7	2.9%	1.5%	0.1%

Table 4.2: Average VOI (left) and VCC (right) with respect to small batch sizes (Q)

4.2.3 Size of the Penalty Costs

The VOI and the VCC also depend on the size of the penalty for mismatches between supply and demand as reflected in the parameters m_0 and m_1 (the retailer's and supplier's product margin), and the supplier's expediting cost b. As the product margin for either facility decreases, VOI increases. We show this relationship in Table 4.2 where the values for VOI and VCC are averaged across experiments at each level of m_0 and m_1.

		VOI				VCC			
Retailer Margin		40%	50%	60%	Mean	40%	50%	60%	Mean
Supplier Margin	40%	5.5%	4.3%	3.7%	4.5%	7.0%	6.4%	6.0%	6.5%
	50%	5.0%	4.1%	3.5%	4.2%	5.6%	5.3%	4.9%	5.2%
	60%	4.5%	3.7%	3.1%	3.8%	4.5%	4.2%	3.8%	4.8%
Mean		5.0%	4.0%	3.4%	4.2%	5.7%	5.3%	4.9%	5.6%

Table 4.2: Sensitivity of the VOI and the VCC to product margin

For the retailer, when the cost of the product is high, the cost of outdating is also high relative to the opportunity cost of a lost sale. Hence, without information sharing, the retailer holds less inventory to avoid costly outdating. Fresher product provided through information sharing reduces the prospect of outdating and enables the retailer to achieve a higher service level that enhances revenues for both the retailer and supplier. Conversely, when the cost of the product is low, the opposite is true and the retailer has a higher service level even *without* information sharing so that *with* information sharing, the major benefit is primarily a reduction in the retailer's outdating. In turn, this negatively impacts the supplier's expected profit. Hence, the opportunity for improving total supply chain profit is greater with a lower retailer margin.

The same relationship holds for the supplier's margin, as lower margins translate into a higher cost of expediting cost for the supplier. This arises because we predicate the expediting cost on the supplier's purchase cost and hence the supplier is more likely to order earlier without information sharing – thereby decreasing the retail shelf life.

5. Extensions

In this section we explore model extensions that include 1) minimum product freshness and supplier price sensitivity to freshness, 2) analysis of the optimal order quantity and its impact on both the VOI and the VCC, and 3) demand sensitivity to product freshness.

5.1 Price Sensitivity to Freshness and Minimum Product Freshness

In our earlier analysis, we assume that the supplier receives the same revenue per unit, regardless of its product freshness, and that the retailer accepts delivery of product without regard to its remaining lifetime. From a practical perspective, however, it is reasonable to expect that 1) a supplier with fresher product may obtain a higher price than a supplier with older product and 2) the retailer may refuse shipment if the remaining product lifetime is too short. Thus, we test how these two relaxations affect the VOI and the VCC.

With regard to supplier pricing, we now assume a simple linear model of freshness dependent pricing where the supplier's revenue per unit is increasing with respect to its product freshness. Let $p_1 = (1 - m_0)$ denote the supplier's maximum revenue per unit. Now let $p_{1,A}$ denote the revenue per unit for inventory at the supplier with a remaining retail shelf life of A days. By definition, we assume that $p_{1,M} = p_1$. Then,

$$p_{1,A} = p_1 - p_1 \delta \left(1 - \frac{A}{M}\right),$$

where $0 \le \delta \le 1$ is a pricing constant that conceptually represents sensitivity to freshness.

With regard to ensuring a minimum level of product freshness for the retailer, we explore the case in which the supplier is restricted from shipping product with less than $A_{\min}$ days of remaining lifetime. We define $A_{\min}$ as the minimum lifetime in which the expected profit from a replenishment of Q units is strictly positive. Now, let $\phi_A(\cdot)$ denote the A-fold convolution of demand and let $\phi_1(\cdot) \equiv \phi(\cdot)$. For $A \ge 2$ we have $\phi_A(x+y) = \sum_x \sum_y \phi(x)\phi_{A-1}(y)$. Then

$$A_{\min} = \min\left(A : \sum_{D=0}^{\infty} \left[-p_{1,A}(Q-D)^+ - h_0 A \left(\frac{Q-(Q-D)^+}{2} \right) + \left(p_0 - p_{1,A}\right)\min(Q,D) \right] \phi_A(D) > 0 \right). \quad (8)$$

On the right side of (8), A is conditioned on the expected cost of product outdating, the approximate expected holding cost, and expected profit contribution. An immediate consequence of the minimum freshness constraint is that inventory may now expire at the supplier. Assuming that the next period marks the β period from the last time the retailer ordered, if the supplier places a replenishment order this period it faces a probability of the product outdating before the next retailer's order of $P(\tilde{D} \geq M - A_{\min} + \beta)$. When it becomes obvious the supplier's inventory will expire the next period, the supplier places a replenishment order so as to avoid the penalty b. We assume that the time between orders is small enough the supplier never incurs an outdating cost for this second replenishment.

Accommodating both minimum product freshness and price dependent freshness for the retailer's replenishment decision in the NIS and DIS cases requires a trivial modification to the formulations expressed in (1) and (4) by replacing the term representing the retailer's purchase cost: i.e., replace $-q\left(1-m_0\right)$ with $-qp_{1,A}$. The supplier's policies, however, are fundamentally different and considerably more complex. Details are provided in Appendix B. For the CC case, the policies are unchanged as the supplier's price is meaningless with centralized control.

With our changed assumptions, we explore the VOI and the VCC in a numerical study of 576 experiments that comprises a factorial design of the following parameters:

$$\delta \in \left(0.0, 0.1, 0.2, 0.4\right) \quad Q \in \left(6, 7, ..., 11\right) \quad C \in \left(0.45, 0.65\right)$$

$$m_0 \in \left(0.4, 0.6\right) \quad m_1 \in \left(0.4, 0.6\right) \quad b \in \left(0.1, 0.2, 0.3\right)$$

The remaining parameters are fixed across experiments where $M = 5$, $\mu = 6$, and the unit holding costs h_0 and h_1 are 40% of the purchase cost measured on an annual basis.

The main results from the study indicate that 1) the VOI and the VCC decrease with respect to δ and 2) that in the DIS case the supplier's share of the total improvement in supply chain profit increases with respect to δ. Sensitivity with respect to the remaining parameters is the same as in the fixed supplier price case. In Table 5.1 we report the average VOI and VCC for each fixed level of δ.

	Supplier Price Sensitivity (δ)							
	VOI				VCC			
Percentile	0.0	0.1	0.2	0.4	0.0	0.1	0.2	0.4
0.00	0.0%	0.0%	0.0%	0.0%	0.1%	0.1%	0.1%	0.1%
0.25	0.2%	0.2%	0.0%	0.0%	0.5%	0.5%	0.6%	0.6%
0.50	0.7%	0.6%	0.1%	0.0%	0.9%	1.0%	1.1%	1.2%
0.75	1.6%	1.3%	0.9%	0.1%	1.8%	1.8%	1.9%	2.2%
1.00	9.8%	5.9%	3.8%	1.2%	10.6%	7.0%	5.8%	4.5%

Table 5.1: VOI and VCC at percentiles for each value of δ

As observed in Table 5.1, both the range and median values of the VOI and the VCC decrease as δ increases. Overall, the VOI and the VCC are much smaller than in the fixed supplier price case, with averages across all experiments of 0.9% and 1.7%, respectively. Only in the experiments with large batch sizes, $Q \in (10,11)$, and small freshness sensitivity, $\delta \leq 0.1$, do we find instances of any substantial value ($\geq 5\%$).

As δ increases, the supplier is increasingly price motivated to sell the freshest product possible in the NIS Case. The prospect of outdating at the supplier also contributes to a fresher product for sale. Hence, while we find that, on average, there is over a 10% improvement in the supplier's product freshness for $\delta = 0.0$ in the DIS case, this measure drops to 1.2% for $\delta = 0.4$.

As for supplier outdating, we only find measurable levels at $Q \in (10,11)$ since at this batch size relative to mean demand, the retailer requires a minimum lifetime of two days and the retailer's order interval can exceed the allowable product lifetime available for sale at the supplier. For these instances, the average level of outdating is 2.2% of the average quantity purchased per period with a maximum of 8.4%. This compares with an average level of outdating of 3.4% for the retailer and a maximum of 8.5%.

The freshness dependent pricing at the supplier also affects the share of value captured by the retailer and the supplier in the DIS Case. As δ increases, the supplier's share increases, albeit of a decreasing total. In Table 5.2 we report the average share of total profit for the retailer and supplier at fixed levels of δ.

Supplier Price Sensitivity (δ)	0.0	0.1	0.2	0.4
% Supplier	-16.7%	55.5%	60.2%	91.0%
% Retailer	116.7%	44.5%	39.8%	9.0%

Table 5.2: % Share of value in the DIS Case for each value of δ

Note in Table 5.2 that values exceeding 100% represent cases where one firm captures all of the value while the other firm is harmed. Hence we see that for $\delta = 0.0$ the supplier is on average worse off with information sharing (matching the results from §4), but as δ increases, the supplier gains an increasing portion of the total value; for $\delta = 0.4$ the supplier gains more than 91% of the total value. This arises because there is little more that the supplier can do with information to increase product freshness (1.2% on average) and hence the only benefit remains with the supplier's ability to reduce its own penalty and holding costs, which are a very small portion of total costs – hence the lower VOI for large δ.

5.2 Assessing the Optimal Order Quantity

So far in our analysis, we assume the batch size Q is exogenously determined. While our model is explicitly designed to explore the VOI and the VCC, we can also use it to find the optimal Q by searching for the largest total supply chain profit over the range of Q for which it is viable to stock and sell the product. We surmise that total profit is concave with respect to Q. Consider Q_{min} and Q_{max} which represents minimum and maximum values for Q in which the product is market viable. Any value less than Q_{min} poses an unacceptable level of service for the retailer and any value greater than Q_{max} poses an unacceptable level of product outdating. As Q increases between Q_{min} and Q_{max}, the service level increases and so does product outdating. Hence, there is an explicit tradeoff between increasing revenue and increasing outdating cost.

We explore this tradeoff using the experiments from §5.1 by evaluating the total supply chain profit in each case for a fixed set of parameter values as Q changes from 6 to 11. In all comparisons, total profit is indeed concave with respect to Q. We illustrate this general relationship for each supply chain structure in Figure 5.1, by taking the average of total profit across all experiments for each value of Q. Over the range of Q studied, the maximum difference in total supply chain profit by choosing a non-optimal value of Q is 10.2%, the average is 3.1% and the minimum is 2.2%. Figure 5.1 also indicates that the optimal value of Q increases with information and centralized control. In the DIS Case, we find that in 13 sets of comparisons (13.3%), the optimal value of Q increases relative to the NIS Case. For the CC Case, in 60 sets of comparisons (61.2%), the optimal value increases relative to the NIS Case.

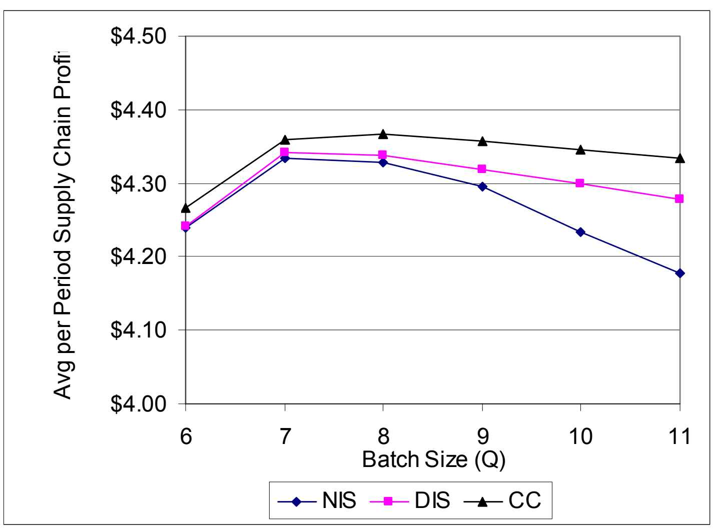

Figure 5.1: Average Total Profit at each value of Q

If we examine the VOI and the VCC in cases where the optimal value of Q is chosen for each supply chain structure (NIS, DIS, CC), then both the VOI and the VCC are minimal. For the DIS Case, the VOI has an average of 0.2% and a maximum of 1.0%. For the CC Case, the VCC has an average of 0.6% and a maximum of 1.8%. Thus, information sharing and centralized control are less valuable if a supply chain can choose the optimal batch size.

5.3 Demand Sensitivity to Product Freshness

In this section, we explore the impact of consumer demand sensitivity to product freshness. A key motivation of our study is that perishables and their quality have together become a key order winning criteria of consumers. Hence, it is reasonable to assume that the rate of demand is sensitive to the freshness of product a retailer makes available for sale. Here, we adopt a simple linear model of demand sensitivity similar to Ferguson and Ketzenberg (2005). Mean demand μ_t in day t is a function of 1) a maximum rate of demand μ, 2) the average lifetime of inventory available for sale at the retailer ω_t relative to the maximum lifetime M, and 3) a constant γ that represents demand sensitivity to product freshness, where $0 \le \gamma \le 1$. Here,

$$\omega_t = \frac{\sum_{x=1}^{M} x i_{x,t}}{I_t} \quad \text{and} \quad \mu_t = \mu - \mu\gamma\left(1 - \frac{\omega_t}{M}\right).$$

We assume that the coefficient of variation in each period t is independent of the mean demand rate so that for any t, total demand D is a random variable with mean μ_t and coefficient of variation C. Note that if $\gamma = 0$, then $\mu_t = \mu$ for all t, corresponding to the case where demand is insensitive to product freshness.

To highlight the impact of demand sensitivity, we choose a set of examples where we know a priori that the VOI and the VCC is low when the mean demand is constant. Here, we fix $\mu = 6$,

$C = 0.5$, $m_0 = m_1 = 0.5$, $M = 4$, and $h = 0.4c$ (measured on an annual basis) and choose a factorial design based on the following parameters:

$$Q \in \{6, 7, ..., 10\},\ b \in \{0.1c, 0.2c\},\ \delta \in \{0.0, 0.2\},\text{ and } \gamma \in \{0.0, 0.1, ..., 0.4\}.$$

When demand is sensitive to product freshness, the VOI and the VCC can be quite substantial. In Table 5.3, we report our summary results for the VOI and VCC across experiments according to percentiles of the set of experiments.

Percentile	Sensitivity to Product Freshness (γ)									
	VOI					VCC				
	0.0	0.1	0.2	0.3	0.4	0.0	0.1	0.2	0.3	0.4
0.00	0.0%	0.0%	0.8%	1.9%	3.3%	0.4%	0.4%	1.1%	2.0%	3.3%
0.25	0.3%	0.7%	3.9%	3.3%	4.3%	0.5%	0.8%	4.0%	3.4%	4.4%
0.50	0.6%	1.3%	5.0%	5.5%	6.4%	0.8%	1.5%	5.1%	5.6%	6.5%
0.75	1.1%	2.3%	5.6%	7.2%	9.8%	1.2%	2.5%	5.7%	7.4%	9.9%
1.00	1.6%	3.1%	7.5%	7.7%	11.1%	1.9%	3.4%	7.6%	7.9%	11.2%

Table 5.3: Summary results of the VOI and the VCC with respect to γ

In Table 5.3, it is clear that both the range and magnitude of the VOI and the VCC increase in γ. With demand sensitivity, information sharing that provides a fresher product provides the capability of increasing the mean demand rate, in addition to reducing product outdating and the service level. Moreover, as γ increases, increasing the demand rate plays an increasingly greater role in the overall profit improvement. For comparison, Table 5.4 shows the average percentage change in outdating, service, and mean demand at fixed levels of γ

Metric	Sensitivity to Product Freshness (γ)									
	VOI					VCC				
	0.0	0.1	0.2	0.3	0.4	0.0	0.1	0.2	0.3	0.4
Outdating	-12.4%	-19.8%	0.4%	-11.8%	-16.4%	35.3%	6.5%	22.0%	3.2%	-14.5%
Service	0.9%	1.0%	3.6%	2.5%	2.7%	1.4%	1.4%	3.9%	2.8%	2.6%
Demand	0.9%	1.6%	4.9%	5.1%	6.4%	1.4%	1.9%	5.0%	5.2%	6.3%

Table 5.4: Average % change in outdating, service, and mean demand with respect to γ

We note that the analysis holds for moderate to large values of δ delta as we illustrate in Table 5.5 with the average VOI for each value of δ and γ.

δ	γ 0.0	0.1	0.2	0.3	0.4
0.0	0.9%	1.6%	4.9%	5.1%	6.4%
0.2	0.3%	1.0%	5.0%	5.6%	6.9%

Table 5.5: The average VOI at fixed levels of δ and γ

6. Conclusion

We separately study the benefits of information sharing and centralized control in a serial supply chain providing a perishable product. Our policies and parameter values are motivated by the offerings of perishable products in the grocery industry. Specifically, we consider a supply chain model consisting of a single product with a fixed lifetime, one supplier, and one retailer that is constrained to order in exogenously determined fixed lot sizes. The retailer orders according to the results of a MDP that balances its holding and spoilage cost with the cost of potential lost sales. Because of the fixed order size, the supplier faces uncertainty only in the timing of the retailer's orders, thus it manages this uncertainty using a safety lead-time that balances its holding cost with the cost of expediting orders that cannot be filled through stock on hand. Finally, the information shared includes the retailer's inventory state (including the age vector of its stock) and its replenishment policy.

We first propose policies for both supply chain members under no information sharing and decentralized control, providing an exact analysis for the expected profits of each firm. Through a numerical study, we compare these results against those obtained by a decentralized control supply chain with information sharing and against a centralized control supply chain where information sharing is implied. We find that supply chains for perishable products benefit the most from information sharing or centralized control when product lifetimes are short, batch sizes are large, demand uncertainty is high, and when the penalty for mismatches in supply and demand are large. The benefits of information sharing alone, however, are not shared equally between the retailer and the supplier. In a decentralized control supply chain, the retailer receives the larger average benefit and in many cases the supplier can be harmed. We show through a model extension, however, that if the supplier's revenue is freshness dependent, the supplier gains a more equitable share, although the VOI in these cases is considerably smaller.

There are a number of important issues still to be addressed. As our numerical test show, an increase in the total supply chain profit is not always Pareto improving for both members. While we look at the value of information sharing and of centralized control, we do not propose contracts that provide firms with the incentive to share/use the information or to act in a centralized manner. Other areas for future research include the modeling of distribution supply chains, longer lead-times, different issuing policies, and capacity restrictions on the supplier.

References

Aviv, Y. and A. Federgruen (1998), "The Operational Benefits of Information Sharing and Vendor Managed Inventory (VMI) Programs," Olin School of Business Working Paper.

Bai, L., (2004), Inventory Control and Demand Distribution Characterization, Unpublished Dissertation from the Georgia Institute of Technology.

Bourland, K. E., S. G. Powell, and D. F. Pyke (1996), "Exploiting Timely Demand Information to Reduce Inventories," *European Journal of Operational Research*, 92, pp. 239-253.

Cachon, G. and M. Fisher (2000), "Supply Chain Inventory management and the Value of Shared Information," *Management Science*, 46(8), 1032-1048.

Chen, F. (1998), "Echelon Reorder Points, Installation Reorder Points, and the Value of Centralized Demand Information," *Management Science*, 44(12), S221–234.

Cohen, Morris A., William P. Pierskalla, and Hsiang-chun Yen (1981), "An Analysis of Ordering and Allocation Policies for Multi-Echelon, Age-Differentiated Inventory Systems," *TIMS Studies in the Management Sciences*, 16, 353-378, Derman, C. and M. Klein (1958), "Inventory Depletion Management," *Management Science*, 4, 450-456.

Feller, W., (1949), "Fluctuation theory of recurrent events," *Transactions of the American Mathematical Society*, 67, 98-119.

Ferguson, M. and M. Ketzenberg (2005), "Sharing Information to Manage Perishables," Working paper Georgia Institute of Technology.

Fries, B. (1975), "Optimal Ordering Policies for Perishable Inventory–II," *Operations Research*, 23(1), 46-61.

Fujiwara, Okitsugu, Hanijanto Soewandi, and Dayani Sederage (1997), "An optimal ordering and issuing policy for a two-stage inventory system for perishable products," *European Journal of Operational Research*, 412-424.

Gavirneni, S., R. Kapuscinski and S. Tayur (1999), "Value of Information in Capacitated Supply Chains," *Management Science*, Vol. 45(1), 16-24.

Goh, C., B. Greenberg, and H. Matsuo (1993), "Two-stage Perishable Inventory Models," *Management Science*, 39(5), 633-649.

Hennessy, T. (1998), "Where Category Management Really Counts," *Progressive Grocer*, 77(2), 103-106.

Kulkarni, V., Modeling and Analysis of Stochastic Systems, Chapman & Hall, 1995.

Kurt Salmon Associates (1993), "Efficient Consumer Response: Enhancing Consumer Value in the Grocery Industry," Kurt Salmon Associates.

Lee, H.L., P. Padmanabhan, and S. Whang (1997), "Information Distortion in a Supply Chain: the Bullwhip Effect," *Management Science*, 43(4), 546-558.

Lee, H., K. So, and C. Tang (2000), "Information Sharing in Supply Chain," *Management Science*, 46(5), 626-643.

Lehmann, E., (1990), Elements of Large Sample Theory. Springer-Verlag, New York.

Leiberman (1958), "LIFO vs. FIFO in Inventory Depletion Management," *Management Science*, 5, 102-105.

Moinzadeh, K., (2002), A Multi–Echelon Inventory System with Information Exchange, *Management Science*, 48(3), 414–426.

Nahmias, S. (1975), "Optimal Ordering Policies for Perishable Inventory–II," *Operations Research*, 23(4), 735-749.

Nahmias, S. (1982), "Perishable Inventory Theory: A Review, *Operations Research*," 30(4), 680-707.

Nahmias, S. and S. Smith (1994), "Optimizing Inventory Levels in a Two Echelon Retailer System with Partial Lost Sales," *Management Science,* 40(5), 582–596.

Office of Technology Assessment Report (1979), "Open Shelf-Life Dating of Food," NTIS order #PB80-101629, August.

Pierskalla and Roach (1972), "Optimal Issuing Policies for Perishable Inventory," *Management Science*, 18(11), 603-614.

Prastacos, Gregory P. (1981), "Allocation of a Perishable Product Inventory," *Operations Research*, 29(1), 95-107.

Putterman, M., Markov Decision Processes: Discrete Stochastic Dynamic Programming, John Wiley & Sons, Inc., 1994.

Raafat, F. (1991), "Survey of Literature on Continuously Deteriorating Inventories*," Journal of the Operational Research Society*, 42(1), 27-37.

Raghunathan, S., (2001), "Information Sharing in a Supply Chain," *Management Science* 47(4), 605-610.

Ross, S., (1996), Stochastic Processes. John Wiley & Sons, New York.

Saporito, B. (1995), "What's for Dinner? The Battle for Stomach Share," *Fortune*, 131(9), 50-57.

Silver, E., Pyke, D., and R. Peterson, (1998), Inventory Management and Production Planning and Scheduling, Wiley.

Yen, H. (1965), "Inventory Management for a Perishable Product Multi-echelon System," Unpublished Ph.D. Dissertation, Department of IE/MS, Northwestern University, Evanston, IL.

Appendix A
Retailer Order Probabilities in the NIS Case

Here, we characterize the probability $P(q' = Q \mid \beta)$ introduced in §3.1.2. Without information sharing, the supplier only knows the batch size Q and the history of the number of periods since the retailer's last order β . We follow the procedure outlined in Bai (2005) to show how this information is used to determine the retailer's order distribution.

Let X_i be a random variable representing the usage of the product (sales and outdating) at the retailer on day i for $i = 1, \ldots, M$. The X_i s are independent with the same mean and variance, but they may come from different distributions. Assuming the retailer uses a reorder point inventory control policy (a reasonable assumption in this industry), once the retailer's approximate inventory position I_i is below the reorder point r, then an order quantity of size Q will be ordered at time t_i. Thus, during the time interval $[t_{i-1}, t_i)$ with length $\tilde{D}_i = t_i - t_{i-1}$, the relationship between accumulated usage and the retailer's inventory position can be expressed as $I_i = I_{i-1} + Q - \sum_{j=1}^{\tilde{D}_i} X_j$. Then the accumulated usage during time interval $\tilde{D}_i$ is $\sum_{j=1}^{\tilde{D}_i} X_j = I_{i-1} + Q - I_i$. Therefore, an interval length $\tilde{D}$ can be defined by the minimal value of n for which the nth accumulated usage is greater than Q, that is,

$$\tilde{D} = N(Q) + 1 \equiv \min\{n : S_n = X_1 + X_2 + \cdots + X_n > Q\}, \qquad \text{(A.1)}$$

where $N(Q) \equiv \max\{n : S_n = X_1 + X_2 + \cdots + X_n \le Q\}$.

The following lemma from Feller (1949) provides the reasoning basis of the first two moments of the demand distribution for deriving the estimates.

LEMMA. *If the random variables* $X_1, X_2, \ldots$ *have finite mean* $\mathrm{E}[X_i] = \mu$ *and variance*

$\mathrm{Var}[X_i] = \sigma^2$, *and* $\tilde{D}$ *is defined by* (A.1)*, then* $\mathrm{E}[X_i]$ *and* $\mathrm{VAR}[X_i]$ *are given by:*

$$\mathrm{E}[\tilde{D}] = \frac{Q}{\mu} + o(1) \quad \textit{and} \quad \mathrm{Var}[\tilde{D}] = \frac{Q\sigma^2}{\mu^3} + o(1) \quad \text{as } Q \to \infty \quad \textit{respectively.}$$

The next theorem provides the asymptotic distribution of $\tilde{D}$. Its proof is a trivial extension to Theorem 3.3.5 in Ross (1996).

THEOREM. *Under the assumptions of the Lemma,* $\tilde{D}$ *has the asymptotic normal distribution with mean* Q/μ *and variance* $Q\sigma^2/\mu^3$:

$$\tilde{D} \to \mathrm{N}(Q/\mu, \sqrt{Q\sigma^2/\mu^3}) \text{ as } Q \to \infty.$$

According to Theorem 2.7.1 of Lehmann (1990), the theorem still holds even when the daily usages are not identically distributed, but are independent with finite third moments. While an asymptotic distribution may cause concern for small values of *Q*, our simulation studies show it provides good estimates for the distribution parameters over the values of *Q* used in this paper.

Let $\Phi_{\tilde{D}}$ represent the cdf of $\tilde{D}$ with a mean of Q/μ and a variance of $Q\sigma^2/\mu^3$. Thus, $\sum_{y=1}^{\beta} P(q' = Q \mid y) = \Phi_{\tilde{D}}(\beta)$. The supplier's safety lead-time is based on its critical fractile, determined from its cost of being early or late with a replenishment order. Finally, recall that α represents the number of periods the supplier waits after receiving an order before it places a replenishment order. As it is impossible by our assumptions on the retailer's policy that α will exceed the product's lifetime, the supplier places his replenishment order $\alpha = \Phi_{\tilde{D}}^{-1}(\frac{h_1 Q}{b + h_1 Q})$ periods after receiving an order from the retailer.

Appendix B
Supplier's Policies with Model Extensions

In §5.1, we introduce two model extensions, namely freshness dependent pricing for the supplier and a minimum level of guaranteed product freshness for the retailer that together fundamentally change the supplier's replenishment problem for the NIS and DIS cases. Here, we characterize these policies.

NIS Case

The supplier's objective is to maximize profit over the time until the next retailer's order. As in our base model, the maximum time between successive retailer orders is *M* days. Let $\psi_{\tilde{D}}(\beta)$ denote the probability of the retailer placing a replenishment order β days after the last order, $\beta \in (1, 2, ..., M)$. The supplier's decision is to choose a value for α so that expected profit is maximized, as expressed by:

$$\max_{\alpha}\left(\sum_{\beta=1}^{M}\begin{cases}\left[Q\left(p_{1,M}-c_1\right)-b\right]\psi_{\tilde{D}}(\beta) & \alpha \geq \beta \\ Q\left[\left(p_{1,M-\beta-\alpha+1}-c_1\right)-h_1\left(\beta-\alpha-1\right)\right]\psi_{\tilde{D}}(\beta) & \alpha < \beta, M-\beta+\alpha+1 \geq A_{\min} \\ Q\left[\left(p_{1,M-\beta+\alpha-A_{\min}+2}-2c_1\right)-h_1\left(\beta-\alpha-1\right)\right]\psi_{\tilde{D}}(\beta) & \alpha < \beta, M-\beta+\alpha+1 < A_{\min}\end{cases}\right). \quad \text{(B.1)}$$

The expectation of the suppliers profit (B.1) is taken over all probabilities for the retailer ordering within the next *M* days and takes into consideration three conditions: 1) $\alpha \geq \beta$, the case when the retailer orders prior to the supplier receiving replenishment so that the retailer's replenishment is satisfied through expediting, 2) $\alpha < \beta$ and $M-\beta+\alpha+1 \geq A_{\min}$, the case where the supplier holds inventory at the time it receives a retailer replenishment order and that no inventory at the supplier has outdated in the previous $\beta-1$ days. In this case, the supplier obtains a price per unit of $p_{1,M-\beta-\alpha+1}$ and incurs holding cost for $\beta-\alpha-1$ days, and

3) $\alpha < \beta$ and $M - \beta + \alpha + 1 < A_{\min}$, the case when the retailer orders after product has outdated at the supplier. Note that in this case, the supplier replenishes two times between successive retailer orders.

It remains to determine $\psi_{\tilde{D}}(\beta)$. Unlike the base model, a complication arises because the supplier's policy may affect the retailer's order probabilities since the purchase cost to the retailer is dependent on product freshness at the supplier. To partially mitigate this problem, we use the following solution procedure. 1) Determine $\psi_{\tilde{D}}(\beta)$ in the same manner as expressed in Appendix A. 2) Solve for the supplier's optimal policy. 3) Solve for the retailer's optimal policy. 4) Resolve for the supplier's optimal policy using the exact order probabilities that result from the analysis of the retailer's steady state behavior arising from step 3. 5) Resolve for the retailer's optimal policy using the supplier's updated policy. Note that this procedure does not guarantee convergence. That is, the order probabilities that arise from step 5) may be different from step 3) and therefore the supplier's optimal policy may be different than what was solved for in step 4. Note that resolving over multiple iterations still does not guarantee convergence.

To assess the impact this may have on our analysis, we took the 576 experiments that we evaluate in §5.1 and compared the solutions from the first and second iterations. We found that in 18% of the experiments, the policies demonstrated differences, but that the impact on expected profit for either facility was less than 5%. From these comparisons, we find that our solution procedure is suitable for the purposes of our analysis.

DIS Case

In this case, the supplier's optimal policy is unknown, but state dependent on the retailer. We formulate the problem as a MDP with the objective to maximize average expected profit per period. The extremal equations are

$$g\left(\vec{i},A\right)+\overline{c}=\max_{\lambda\in(0,1)}\left(-\lambda c_1+\begin{cases}\left(p_{1,M}-\dfrac{c_1-b}{Q}\right)q^*_{\vec{i},A}+\sum_{D=0}^{\infty}g\left(\tau\left(\vec{i},D,q^*_{\vec{i},A},M\right),A'\right)\phi(D) & A=0\\ p_{1,A}q^*_{\vec{i},A}-hQ+\sum_{D=0}^{\infty}g\left(\tau\left(\vec{i},D,q^*_{\vec{i},A},A\right),A'\right)\phi(D) & A>0\end{cases}\right).$$

As in §3.2.2, the retailer and supplier replenishment decisions are inter-related and decision–making is decentralized. Hence we solve $f\left(\vec{i},A\right)$ for the retailer and $g\left(\vec{i},A\right)$ for the supplier simultaneously and use the same solution procedure for determining q^* and λ^* as expressed in §3.2.2 for the base model.

Appendix C
Detailed Sensitivity Analysis

				Performance Measures in the DIS Case Relative to the NIS Case*					
				Retailer					
Parameter	Value	VOI	VCC	Service	Outdating	Order Quantity	Order Interval	Freshness	Supplier Freshness
Coefficient of Variation	0.5	2.7%	3.3%	1.7%	-34.1%	-0.8%	0.9%	16.5%	20.2%
	0.6	2.8%	3.7%	1.9%	-18.4%	-0.2%	0.4%	14.8%	19.0%
	0.7	7.0%	8.9%	6.1%	-4.6%	4.2%	-3.6%	14.4%	21.5%
Expediting Cost	0.05	4.6%	5.4%	3.3%	-37.7%	0.3%	-0.1%	19.2%	23.8%
	0.10	4.2%	5.4%	3.2%	-21.8%	0.9%	-0.6%	15.8%	20.6%
	0.15	4.0%	5.3%	3.2%	-11.1%	1.4%	-1.1%	13.7%	18.9%
	0.20	3.9%	5.1%	3.2%	-5.4%	1.7%	-1.4%	12.4%	17.7%
Product Lifetime	5	5.8%	8.2%	4.2%	-15.4%	0.8%	-0.4%	18.8%	29.3%
	6	4.2%	4.9%	3.3%	-20.3%	1.0%	-0.7%	16.9%	19.4%
	7	2.4%	2.7%	2.2%	-21.4%	1.4%	-1.3%	10.0%	12.1%
Supplier Margin	0.4	4.5%	5.7%	3.3%	-18.8%	1.1%	-0.8%	15.3%	20.3%
	0.5	4.2%	5.3%	3.3%	-18.9%	1.1%	-0.8%	15.3%	20.3%
	0.6	3.8%	4.9%	3.1%	-19.4%	1.0%	-0.7%	15.3%	20.2%
Retailer Margin	0.4	5.0%	6.5%	3.6%	-18.4%	1.2%	-0.9%	17.4%	21.6%
	0.5	4.0%	5.2%	3.2%	-18.9%	1.2%	-0.9%	14.6%	19.9%
	0.6	3.4%	4.2%	2.8%	-19.7%	0.8%	-0.5%	13.8%	19.2%
Batch Size	8	2.1%	3.3%	2.1%	-3.1%	2.1%	-1.8%	8.5%	10.6%
	9	4.8%	6.0%	3.6%	-21.2%	1.0%	-0.6%	18.8%	21.3%
	10	5.6%	6.5%	3.9%	-32.8%	0.1%	0.1%	18.5%	28.9%

* Performance measures in the DIS Case are calculated as the % change of the measure in the NIS Case. All measures are per period averages, computed from steady state behavior of the MDP. Freshness is measured as the average remaining lifetime at the point of sale.