Information Sharing to Improve Retail Product Freshness of Perishables

Mark Ferguson
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

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

August 18, 2004

Information Sharing to Improve Retail Product Freshness of Perishables

Abstract

We explore the value of information (VOI) in the context of a retailer that provides a perishable product to consumers and receives replenishment from a single supplier. We assume a periodic review model with stochastic demand, lost sales, and order quantity restrictions. The product lifetime is fixed and deterministic once received by the retailer, although the age of replenishment provided by the supplier varies stochastically over time.

Since the product is perishable, any unsold inventory remaining after the lifetime elapses must be discarded (outdated). Without the supplier explicitly informing the retailer of the product age, the age remains unknown until receipt. With information sharing, the retailer is informed of the product age prior to placing an order and hence can utilize this information in its decision–making. We formulate the respective scenarios as Markov Decision Processes (MDPs) and measure the VOI as the marginal improvement in cost that the retailer achieves with information sharing, relative to the case when no information is shared. We establish the importance of information sharing and identify the conditions under which relatively substantial benefits can be realized.

.

Keywords: value of information, inventory management, perishable product

1. Introduction

We place our research in the context of food and agribusinesses and specifically, in 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 or up to $150 billion a year, but also subject the firms to losses of up to 15 percent due to damage and spoilage. Further, perishables have become the order winning criteria of consumers, becoming the core reason many consumers choose one supermarket over another (Heller, 2002). These are all powerful incentives for investment in information enabling technologies for the management of perishables. Indeed many suppliers are embarking on supply chain initiatives premised on information technologies. For example,

> Del Monte is focusing on making the retailer's life easier by taking on more of the work through supply partnerships… Technology has been the key to Del Monte's strategy – along with a sophisticated partnering package. Del Monte is working with retailers on accounting, packaging, merchandising, and sales – shared technology that allows broader and richer enhancement of information. (Hennessy, 2000, p. 74)

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. While our research focus is on groceries, the management of perishable inventories is an important problem confronting many other industries including blood banks, food service, pharmaceuticals, chemicals, and increasingly, biotechnology. Yet the grocery industry is particularly appropriate, given current practitioner activity and industry initiatives. In this paper, we introduce a model that extends the research on perishable inventory systems by evaluating a system where the age of the replenishment is uncertain, the retailer orders in batches, and unmet demand results in lost sales:

three highly significant and relevant aspects to the management of perishables in the grocery industry.

While the importance of managing perishables is growing, there has also been a growing interest in the value of information sharing for supply chain management (VOI) as exemplified by recent contributions to the academic literature by Cachon and Fisher. (2000), Lee et al. (2000), and Moinzadeh (2002). Most of the current research has focused on the potential benefits of sharing *downstream* facility information with upstream facilities located closer to the originating suppliers. Such information may include demands characteristics and inventory positions of the downstream firms. The upstream suppliers can then incorporate this information into their decision making process to better match supply with demand. In contrast, the potential benefits with respect to the reverse flow of information (supplier to the retailer) have received scant attention in the literature. We note that both Chen (2002) and Huang et al. (2003) remark on the need for future research in this area. In this respect, we extend the literature on the VOI sharing in this important direction.

We address the VOI in the context of a retailer that provides a perishable product to consumers. Demand is stochastic and unsatisfied demands are lost. The retailer receives replenishments from a single supplier and there is a batch ordering constraint on the ordering decisions. The product lifetime is fixed and deterministic once received by the retailer, although the age of replenishment varies stochastically over time. These assumptions correspond to the widespread use of packaging highly perishable products with expiration dates. Without the supplier explicitly informing the retailer of the product age, the age of any replenishment remains unknown until receipt. Since the product is perishable, any unsold inventory remaining after the lifetime elapses must be discarded (outdated). With information sharing, the retailer is informed of the product age, prior to placing an order, and hence can utilize this information in its decision–making. We formulate the respective scenarios as Markov Decision Processes (MDPs) and measure the VOI as the marginal improvement in profit that a retailer achieves with information sharing, relative to the case when no information is shared. Given the complexity and computational limitations of these policies, we also introduce and test heuristic policies for both the information sharing and no information sharing cases. These heuristic policies perform very well and allow us to extend our analysis of the VOI over a wider range of parameter values.

Through a numerical study, we establish the importance of information sharing and identify the conditions under which relatively substantial benefits can be realized. We find the average VOI to be 9.6% with a range from 0% to 225%. In a sensitivity analysis, we find that the key drivers behind the VOI are: the expected product lifetime, the issuing policy (FIFO or LIFO), and the variability of both the demand and the remaining lifetime of the product upon receipt. We also investigate the impact of information sharing on the supplier and the whole supply chain. The rest of the paper is organized as follows. §2 reviews the literature, §3 defines the model, §4 provides bounds on the optimal policies, §5 presents a numerical study, §6 introduces and tests heuristic policies, and §7 concludes the paper with future research directions.

2. Literature Review

Our research draws on two separate research streams: the literature on 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

There are two problems addressed by the literature on fixed lifetime perishable inventory theory that are related to the problem we address. These include determining reasonable and appropriate methods for issuing inventory and for replenishing 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 Leiberman (1958) and Pierskalla and Roach (1972) address the conditions where issuing the oldest items first (FIFO) and youngest items first (LIFO) are optimal. With constant product utility until outdating, as is the case with our research, FIFO issuing is optimal. Even so, we also address LIFO inventory issuing since it is clear from practice that inventory issuing is not necessarily controllable by a retailer.

Significant research has been done to derive and evaluate replenishment policies for items with a fixed lifetime. Simultaneously, yet independently, Nahmias (1975) and Fries (1975) were the first to derive and evaluate optimal policies for perishable products with lifetimes greater than two periods. In their models, the quantity of product to be outdated is expressed recursively in terms of previous outdates and demands. They formulate their respective problems as cost–minimizing dynamic programs that include both outdating and shortage costs. In both cases, the optimal ordering policy is shown to be non–stationary and dependent on the age distribution of inventory. Unlike our model, the product is assumed to be fresh on receipt (i.e. the amount of remaining lifetime upon receipt does not vary from one replenishment to the next).

Given the multidimensional state of inventory, computation of optimal solutions using dynamic programs on long lifetime products is impractical since the state space expands exponentially with the number of possible age categories. Hence, much of the more recent work has focused on well performing heuristic policies. More recently, Nandakumar and Morton (1993) and Chui (1995) provide approximations for continuous review perishable systems. We also introduce well performing heuristics that are designed to evaluate the VOI in a periodic review system where 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.

2.2 Value of Information

Recently, a few articles have emerged that provide literature reviews and taxonomies that address the VOI for supply chain management. Sahin and Robinson (2002) and Huang et al. (2003) are representative examples, each providing a very broad overview of the literature and offering classification schemes. Chen (2002) is notable for the depth of analysis in exploring the different types of information sharing and explaining and comparing the analytical results among several key contributions to the field. Collectively, the reviews indicate that a preponderance of the research in this area focuses on the value of demand information to improve supply chain performance. Representative examples include contributions by Bourland et al. (1996), Cachon and Fisher (2000), and Moinzadeh (2002). Only a few studies have addressed the value of supply information. For example, Chen and Yu (2002) consider the case where lead-time information is shared forward in the supply chain so that customers can reduce supply uncertainty.

Beyond our own study, Ketzenberg and Ferguson (2003) 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 in a serial supply chain consisting of a single retailer and a single supplier. Here, information is shared upstream, where the retailer shares its age–dependent inventory state, replenishment policy, and demand information with the supplier.

While we also address the value of information with respect to the supply of a perishable product, we examine the reverse flow of information in which the supplier shares its inventory state with the retailer. Also, Ketzenberg and Ferguson (2003) model supply chains where the supplier's ordering policy is highly dependent on the retailer's actions, we model supply chain structures where the supplier provides for a large number of retailers. Thus, the replenishment actions of a specific retailer are considered inconsequential to the choice of ordering policy of the supplier. This type of scenario is more appropriate for the grocery industry.

3. Model

The general setting is a retailer that provides a perishable product to consumers and receives replenishments from a larger supplier. We assume a periodic review inventory model, as this is the most common system used in the grocery industry. The product is perishable and has a maximum retail product shelf life of M periods, although the remaining shelf life at the time of replenishment varies between 1 and M as we later discuss. 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.

The order of events each day follows the sequence: 1) receive delivery, 2) outdate inventory, 3) observe and satisfy demand, and 4) place replenishment order. Retail demand is discrete, stochastic, and stationary over time. Let D denote total demand in the current period, with probability mass function $\phi(\cdot)$, a cdf of $\Phi(\cdot)$, a mean of μ_D, and a coefficient of variation of C_D. Unsatisfied demand is lost. Let p be the selling price and w the per unit purchase cost from the supplier. We assume that the only penalty for a lost sale is the lost margin, p-w. A holding cost h is assessed on ending inventory.

The replenishment decision q is restricted to multiples of a batch quantity Q such that $q = nQ$ in the current period, where $n = 0, 1, 2, \ldots$. The batch quantity Q is given and fixed. This assumption captures certain economies of scale in transportation, handling, or packaging, although we do not model these economies explicitly (i.e., there is no fixed order cost). Such an assumption is common in practice and the literature (see Chen 1998, Cachon and Fisher 2000, Moinzadeh 2002). Although Q is exogenous in our model, we nevertheless evaluate the impact of this important parameter in our analysis.

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 $x = 1, \ldots, M$. Let $\boldsymbol{i} = (i_1, i_2, \ldots, i_M)$ represent the vector of inventory held at each age class and define $I = \sum_{x=1}^{M} i_x$.

We separately explore both FIFO and LIFO inventory issuing policies used to satisfy demand. While it is clear that FIFO issuing is optimal, generally retailers do not have explicit control of how demand is satisfied. Exceptions exist however, such as the load-from-the-back shelving systems often used for dairy products. When control is left to customers, they are apt to select the freshest products first.

Product ordered in period t arrives in period t+1, where t denotes the current period. This is consistent with our discussions with managers in the grocery industry because delivery trucks from the supplier operate daily and may include a stop with as little as one days notice. The retailer orders from a completely reliable exogenous supplier. That is, the supplier has ample capacity so that all retail orders are fully satisfied one period later. The remaining lifetime of replenishment received in any period is a discrete random variable, although all units have the same remaining lifetime. We do not model the supplier explicitly, but rather address the stochastic nature of the product age at the time of replenishment. These assumptions represent the case of a large supplier that provides product to many independently controlled (in terms of ordering policies) retailers, a scenario that is common in the grocery industry today. Since the retailer makes up only a small portion of the supplier's total order quantity in any given period, the supplier's inventory policy is assumed to be independent of the retailer's policy.

Let $A,\ A \in \{1, 2, ..., M\}$, denote the remaining lifetime of replenishment associated with an order placed in the current period. Further, let $\psi(\cdot)$ denote its probability mass function with mean μ_A and coefficient of variation C_A. Without information sharing, A is unknown at the time an order is placed, although the retailer does know $\psi(\cdot)$. Note that this scenario represents common practice in the grocery industry. Retailers do not typically know the age of replenishment until it is received, although they can estimate the age distribution from their replenishment history. We formulate the replenishment problem as an infinite-horizon dynamic program where the objective is to find the retailer's optimal reorder policy so that its expected cost is minimized. 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 realization A of the remaining lifetime for any replenishment.

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. Letting $\boldsymbol{i}'$ denote the retailer's inventory level in the next period and $\tau(\boldsymbol{i}, D, nQ, A)$ denote the one period transfer function, then $\boldsymbol{i}' = \tau(\boldsymbol{i}, D, nQ, A)$ where

$$i'_x = \begin{cases} \left(i_{x+1} - \left(D - \sum_{z=1}^{x} i_z \right)^+ \right)^+ & x \neq A \\ \left(i_{x+1} - \left(D - \sum_{z=1}^{x} i_z \right)^+ \right)^+ + nQ & x = A \end{cases} \quad \text{for FIFO inventory issuing}$$

and

$$i'_x = \begin{cases} \left(i_{x+1} - \left(D - \sum_{z=x+2}^{M} i_z \right)^+ \right)^+ & x \neq A \\ \left(i_{x+1} - \left(D - \sum_{z=x+2}^{M} i_z \right)^+ \right)^+ + nQ & x = A \end{cases} \quad \text{for LIFO inventory issuing.}$$

Given a starting vector of inventory $\boldsymbol{i}$ and an order quantity multiple of n, the infinite horizon cost-to-go, if future periods behave optimally, is $g(n,\boldsymbol{i})$. The order quantity multiple that minimizes the cost-to-go is denoted by $n^*(\boldsymbol{i})$. We represent the one period holding and penalty cost by $L(I)$. Formally:

$$L(I) = h\sum_{D=0}^{I}(I-D)\phi(D) + (p-w)\sum_{D=I}^{\infty}(D-I)\phi(D)\,. \tag{1}$$

We can explicitly write the infinite horizon recursion as:

$$g(n,\boldsymbol{i}) = L(I+nQ) + w\sum_{D=0}^{i_1}\left(i_1 - D\right)\phi(D) + \sum_{D=0}^{\infty}\sum_{A=1}^{M} f\left(\tau\left(\boldsymbol{i},D,nQ,A\right)\right)\psi\left(A\right)\phi(D) \tag{2}$$

where

$$f(\boldsymbol{i}) = \min_{n\geq 0} g(n,\boldsymbol{i})\,. \tag{3}$$

The right hand side of equation (2) computes the total expected cost that is composed of the one-period holding and penalty cost, the cost of any unsold product that perishes this period, and the future expected cost. Note that the expectation of current period costs are predicated only on $\phi(\cdot)$, while the expectation of future cost is predicated on both $\phi(\cdot)$ and $\psi(\cdot)$. The initial state space for n is the set of positive integer values. This formulation is similar to the approaches followed by Fries (1975) and Nandakumar and Morton (1993).

4. Bounds and Policies

In this section, we derive lower and upper bounds on the order quantity to reduce the state space of our optimization problem. We then present finite-horizon versions of equation (2) that provide optimal solutions for the no information sharing and information sharing cases.

4.1 Lower bound

For a lower bound on the order quantity, we assume that all replenishments arrive with a lifetime of only one period, i.e. $\boldsymbol{i} = i_1$. From a perishable product standpoint, this presents a worst case for the retailer and raises its overage cost significantly as all unsold inventory perishes and can not be carried over to future periods. In this case, the retailer's problem reduces to a series of single period newsvendor problems and equations (1), (2) and (3) reduce to:

$$L(n) = w\sum_{D=0}^{nQ}(nQ-D)\phi(D) + (p-w)\sum_{D=nQ}^{\infty}(D-nQ)\phi(D)\,,$$

$$g(n) = L(n) \quad \text{and}$$

$$f(n) = \min_{n\geq 0} g(n)\,,$$

resulting in the classic newsvendor solution: $\underline{q} = \underline{n}Q \geq \Phi^{-1}\left(\dfrac{p-w}{p}\right)$.

The underbar notation on q and n represents the lower bound solution where $\underline{n}$ is the smallest integer that satisfies the inequality condition.

4.2 Upper bound

For an upper bound on the order quantity, we use the solution to the basic infinite-horizon periodic review system for a non-perishable product. It is intuitive that the possibility of a product perishing before it can be sold raises the overage cost for the retailer. Since the underage cost remains the same, the order-up-to level is smaller for a perishable product compared to a non-perishable product. If no outdating occurs, the infinite horizon recursion (see Zipkin, 2000 pg 382 – 385) becomes:

$$g(n,I) = L(I+nQ) + \sum_{D=0}^{\infty} f\left(\tau\left(I,D,nQ\right)\right)\phi(D), \tag{4}$$

where (4) is simply (2) minus the outdating term and (1) and (3) are the same as in the original case. The recursion is minimized when, at the end of each period, the retailer chooses the smallest value of *n* such that $I + \overline{n}Q \geq \Phi^{-1}\left(\frac{p-w}{p-w+h}\right)$, giving an upper bound of $\overline{q} = \overline{n}Q$. The upperbar notation represents the upper bound solution and $\overline{n}$ is the smallest integer that satisfies the inequality condition. This result is proven in Nahmias (1975). We now define our new state space as *N*, consisting of all integer values in the range $\left[\underline{n}, \overline{n}\right]$.

4.3 Finite Horizon Policies

Given the reduced state space provided by the upper and lower bounds, we proceed to modify equation (2) for the finite horizon case of no information sharing. Since all inventory, either on hand or on order, expires within the next *M* periods (i.e. any decision made this period has no effect on the retailer's cost in period *M* + 1), our problem reduces to an *M*-period, finite-horizon MDP. By assuming that the terminal cost is $f(\boldsymbol{i})^{M+1} = -wI$, or equivalently, the retailer may sell back any remaining stock at its purchase cost, the infinite horizon problem is easily reduced to a finite horizon problem. Let the superscript *NIS* denote the no information sharing case. The *M*-period, finite-horizon recursion is

$$g^{NIS}\left(n,\boldsymbol{i}\right)^{t}=\left\{L(I+nQ)+\sum_{D=0}^{\infty}\left[\begin{array}{l}w\left(i_{1}-D\right)^{+}+\\ \sum_{A=1}^{M}f^{NIS}\left(\tau\left(\boldsymbol{i},D,nQ,A\right)\right)^{t+1}\psi\left(A\right)\end{array}\right]\phi\left(D\right)\right\} \tag{5}$$

where $f^{NIS}(\boldsymbol{i})^{t}=\min_{n\in N}g(n,\boldsymbol{i})^{t}$.

Since the state and decision spaces are discrete and finite and the cost is bounded, there exists an optimal stationary policy that does not randomize (Putterman, 1994 pg 102 - 111). Let $n^{NIS}(\boldsymbol{i})^{1}$ denote the optimal stationary policy of order quantity multiples. The resulting optimal cost-to-go is $f^{NIS}(\boldsymbol{i})$, for all $t\in 0,\ldots,M$.

With information sharing, the retailer knows A prior to placing an order in the current period. In this case, the state space is expanded to include this information. Let the superscript *IS* denote the information sharing case. The *M*-period, finite-horizon recursion is:

$$g^{IS}\left(n,\boldsymbol{i},A\right)^{t}=\left\{L(I+nQ)+\sum_{D=0}^{\infty}\left[\begin{array}{l}w\left(i_{1}-D\right)^{+}+\\ \sum_{A'=1}^{M}f^{IS}\left(\tau\left(\boldsymbol{i},D,nQ,A\right)^{t+1},A'\right)\psi\left(A'\right)\end{array}\right]\phi\left(D\right)\right\} \tag{6}$$

where $f^{IS}(\boldsymbol{i},A)^{t}=\min_{n\in N}g^{IS}(n,\boldsymbol{i},A)^{t}$.

Let $n^{IS}(\boldsymbol{i})^{1}$ denote the optimal stationary policy of order quantity multiples. The resulting optimal cost-to-go is $f^{IS}(\boldsymbol{i},A)$, for all $t\in 0,\ldots,M$. Note that while A is known with respect to any order placed in the current period, this information is not known for subsequent periods. Hence, the state transition probability from state $\left(\boldsymbol{i},A\right)$ to state $\left(\boldsymbol{i}',A'\right)$ is predicated on both $\phi\left(\cdot\right)$ and $\psi\left(\cdot\right)$ just as it is in the optimization of equation (5) for the no information sharing case.

Since expected profit is a more appropriate metric for the grocery industry, we interpret the VOI in terms of a change in expected profit due to information sharing by a simple conversion of our cost minimizing policies. Our switch to a profit maximization problem is simplified by the fact that we set the cost of a lost sale equal to the lost margin. Thus, the optimal ordering quantity multiples $n^{NIS}(\boldsymbol{i})^1$ and $n^{IS}(\boldsymbol{i})^1$ are equivalent for both the cost minimization and profit maximization problems. Letting $\pi^{NIS}(\boldsymbol{i})$ and $\pi^{IS}(\boldsymbol{i}, A)$ represent the optimal expected profits from the stationary policy for the no information sharing and information sharing cases respectively, we have:

$$\pi^{NIS}(\boldsymbol{i}) = p\mu_D - f^{NIS}(\boldsymbol{i}) \quad \text{and} \quad \pi^{IS}(\boldsymbol{i}, A) = p\mu_D - f^{IS}(\boldsymbol{i}, A) \ .$$

In the next section, we explore the VOI through an extensive numerical study. 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 (1995, p. 124).

5. Numerical Study

We evaluate the VOI, measured as the % improvement in expected retailer profit, relative to the case where information is not shared. Specifically, define

$$VOI = \frac{\left(\pi^{NIS}(\boldsymbol{i}) - \pi^{IS}(\boldsymbol{i}, A)\right)}{\pi^{NIS}(\boldsymbol{i})}.$$

Consumer demand $\phi(\cdot)$ corresponds to a truncated negative binomial distribution with a maximum value of 50 (any probabilities for demand values 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.

The maximum product lifetime M is 5 periods, although the age of receipts A varies according to $\psi(\cdot)$. The distribution $\psi(\cdot)$ is derived randomly using Excel Solver as follows: The allowable range for A are integers between 1 and M, and we set the probabilities for each integer to achieve the desired μ_A and C_A. There is not a unique distribution that satisfies these constraints, but through a series of numerical tests, we have observed that the shape of the distribution has no measurable impact on the results. This method is analogous to the approach previously tested and used in Souza, et al. (2002).

Each period represents a day and the holding cost is 25% of the purchase cost, measured on an annual basis. We consider a set of experiments that comprise a factorial design for all combinations of the following parameters:

$$\mu_D \in \{3,4\} \quad C_D \in \{0.60, 0.75, 0.90\} \quad Q \in \{1,2,4,6\}$$

$$\mu_A \in \{2,3,4\} \quad C_A \in \{0.2, 0.3, 0.4\} \quad w \in \{0.4, 0.55, 0.70\}$$

Our selection of parameter values reflect a broad range of situations for many short life-time products that include deli items, fresh cut produce, packaged meats and seafood, as well as many packaged produce items (Raper, 2003 and Pfankuch, 2004). These include low and high demand variability, small and large batch sizes (relative to mean demand), low and high variability in the age of replenishment, and low and high-priced inventory. While the mean demand rates are low, they are reasonable, particularly given the ever-increasing number of products over which demand is spread (Kollars, 2001). 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.

We duplicate the full set of 648 experiments contained in our factorial design for the two inventory issuing policies we explore: LIFO and FIFO. Hence, there are a total of 1,296 experiments with which we explore the VOI. In §5.1, we discuss our general observations, in §5.2 we report the results of our sensitivity analysis, and in §5.3 we evaluate the impact of information sharing on the supplier.

5.1 Computational Results and General Observations

In general, we find that the sharing of supply information enables a retailer to purchase considerably fresher product and consequently, information can be quite valuable. In Table 5.1 we report the VOI at given percentiles of the set of 1,296 experiments evaluated. For example, the 0.50 percentile denotes the median values for VOI. Across experiments, the range of the VOI is between 0% and 225%, with a mean of 9.6% and a median of 6.1%. We also report additional performance measures of interest that include the percentage change in the average remaining product lifetime of receipts, remaining product lifetime at the point of sale, level of outdating, and service (fill-rate), where all measures are relative to the no information sharing case. Note that the values for *each* performance measure are ranked according to the percentile (from lowest to highest) and not according to the VOI.

		% Change in			
Percentile	VOI	Lifetime of Receipts	Lifetime of Sales	Oudating	Service
0.00	0.0%	0.4%	-14.7%	-83.8%	-32.8%
0.05	1.4%	3.9%	-4.7%	-61.0%	-4.4%
0.10	1.8%	4.7%	-3.3%	-53.1%	-1.6%
0.25	3.4%	6.6%	-0.7%	-37.8%	-0.1%
0.50	6.1%	11.9%	1.4%	-22.6%	0.8%
0.75	11.6%	17.6%	3.6%	-11.6%	2.5%
0.90	20.5%	22.2%	6.1%	-2.3%	5.5%
0.95	26.9%	24.7%	7.9%	0.0%	7.6%
1.00	224.6%	36.1%	18.2%	0.0%	32.5%

Table 5.1: Performance measures at percentiles of the 1,296 experiments

Although it is clear from Table 5.1 that the VOI can at times be large, the range that is reported also reveals that any realization of value is sensitive to model parameterization. Below, we first discuss the drivers of value and follow with a sensitivity analysis to understand the conditions in which information sharing is most beneficial.

In general, we find that a retailer, by using the supplier's age of inventory in its replenishment decision, will increase the expected lifetime of its replenishment by ordering more in periods when the supplier has fresher product and less in periods when the supplier has older product. On average, the expected improvement in replenishment lifetime increases from 3.0 days to 3.4 days (13.3%). In turn, the level of outdating that arises from product expiration decreases from an average of 0.48 units per period to 0.37 (-22.9%).

The improvement in product freshness is not necessarily shared with consumers. Although on average consumers realize a 1.4% improvement in the remaining lifetime at the point of sale, it ranges from -14.7% to 18.2% as shown in Table 5.1. We find that the change in product freshness to consumers is largely a function of a change in retailer service. When the service level increases, average inventory levels also increase so that product freshness decreases at the point of sale. On average, the retailer observes a slight improvement in the expected service fill rate (1.0%) since the expected cost of over–stocking, relative to the opportunity cost of a lost sale, is reduced with a fresher product. Yet this is not always the case as shown in Table 5.1 where we observe that the service fill-rate actually decreases in approximately 25% of the experiments. We find that information sharing enables a systematic tradeoff between a decrease in the cost of outdating and an increase in profit contribution due to higher service levels so that at times service, is sacrificed for a very significant savings in outdating. On average, we observe that 82% of the improvement in average expected profit arises from a reduction in outdating and 18% arises from higher service. In Figure 5.1, we present the average change in outdating cost and the average change in profit contribution at fixed intervals of the VOI. Note that the outdating cost is reported in absolute values for ease of comparison.

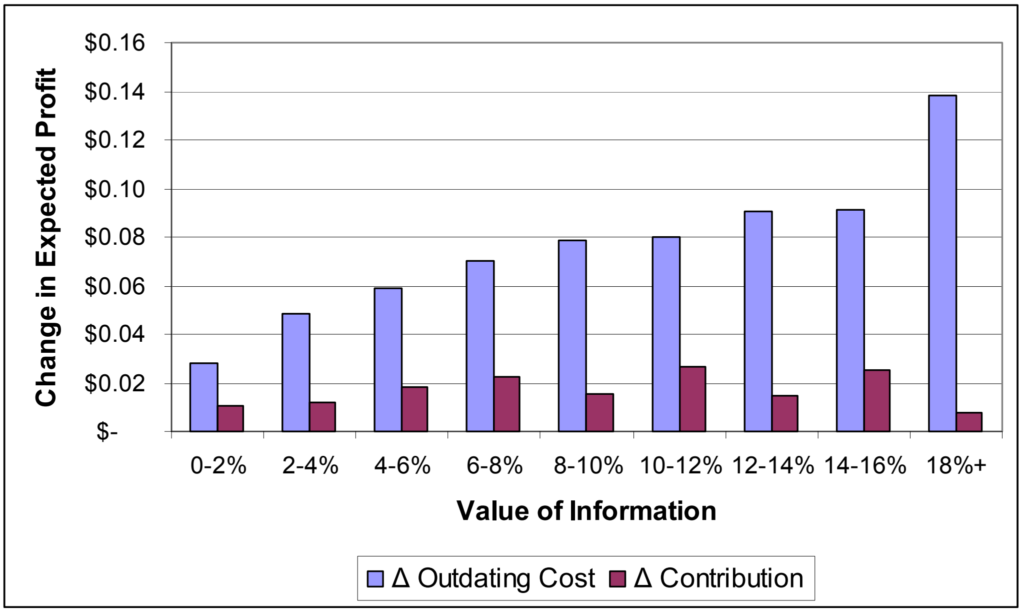

Figure 5.1: The Components of the VOI

As Figure 5.1 illustrates, a reduction in the cost of outdating is largely responsible for the VOI that we observe across experiments. Generally, the improvement in outdating is accompanied by a much smaller improvement in service, although we do find that the increase in profit contribution exceeds the reduction in outdating cost in approximately 10% of the experiments.   We elaborate below with a sensitivity analysis.

5.2 Sensitivity Analysis

In Table 5.2, we report the average VOI across all 1,296 experiments for each fixed parameter value. In addition, we report the associated average change in the cost of outdating, the average change in profit contribution, and the average change in the retailer's order size. Corroborating our prior results, we find the VOI is largely a function of a decrease in the cost of outdating and there is virtually no sensitivity or pattern that we can observe with respect to changes in profit contribution that arise from a change in service fill-rate. The VOI itself is largely a function of product perishability, the ability of the retailer to match supply with demand, and the drivers of retail profitability as we explain below.

Parameter	Value	VOI	Δ Contribution	Δ Outdating Cost	Δ Order Size
Issue Policy	FIFO	10.4%	$0.01	$0.07	-0.10
	LIFO	8.9%	$0.01	$0.05	-0.06
Mean Demand	3	10.8%	$0.01	$0.06	-0.07
	4	8.5%	$0.02	$0.07	-0.09
Demand CV	.60	7.5%	$0.01	$0.06	-0.09
	.75	9.8%	$0.02	$0.06	-0.07
	.90	11.6%	$0.02	$0.07	-.0.08
Expected Lifetime	2	11.8%	$0.01	$0.08	-0.08
	3	9.8%	$0.02	$0.06	-0.05
	4	7.4%	$0.01	$0.05	-0.11
Age CV	.2	3.7%	$0.01	$0.03	-0.04
	.3	9.0%	$0.01	$0.06	-0.07
	.4	16.2%	$0.02	$0.10	-0.13
Product Cost	0.40	3.9%	$0.01	$0.05	-0.09
	0.55	7.4%	$0.01	$0.06	-0.08
	0.70	17.7%	$0.01	$0.07	-0.06
Batch Size	1	7.6%	$0.02	$0.05	-0.04
	2	7.6%	$0.02	$0.05	-0.03
	4	8.7%	$0.02	$0.06	-0.07
	6	14.7%	$0.00	$0.09	-0.18

Table 5.2: Sensitivity of the VOI to Parameters

5.2.1 Product Perishability

As shown in Table 5.2, the VOI is larger as the expected lifetime of the product at replenishment μ_A decreases. When the product lifetime is short, improvements in product freshness have a larger impact than when the product lifetime is long. Improvements in product freshness reduce the potential for outdating, allowing the retailer to carry more inventory for the same amount (or less) of product outdating.

In contrast, long product lifetimes result in smaller VOI 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 information sharing has no effect on retailer behavior because product freshness is no longer material to the problem. Consequently, the VOI is zero.

The issuing policy that is used to satisfy demand is another aspect of product perishability that impacts retailer behavior and hence the VOI. We find that the VOI is higher with FIFO issuing. In Table 5.3 we report the average expected values for retailer profit across experiments for both issuing policies and each level of information sharing.

	LIFO	FIFO	% Change
Without Information	\$1.01	\$1.15	13.9%
With Information Sharing	\$1.10	\$1.27	15.5%
% Change	8.9%	10.4%	

Table 5.3: Comparison of LIFO versus FIFO Issuing Policies

Since the impact of perishability is greater with LIFO issuing, we would expect that the VOI would also be greater with LIFO issuing. Although we do find that in 35% of the cases the VOI is indeed greater, the results in Table 5.3 and in the majority of experiments show that the VOI can be substantially larger with FIFO issuing. We find that the difference is greatest at low values for C_A, C_D, μ_A, and higher values for μ_D.

To better understand these results, it is necessary to dig deeper into the underlying factors that drive profitability with each policy. With FIFO issuing, the retailer has better control of its inventory so that product outdating is minimized and it is profitable to maintain an 85% service fill rate without information sharing. With information sharing, the freshness of replenishment increases by 14.1% and consumers realize a corresponding increase of 2.2% in product freshness at the point of sale. The key is that in any replenishment period, an improvement in the freshness of replenishments decreases the likelihood of product outdating in future periods.

Now with LIFO issuing, the retailer has inherently less control of product outdating so that the cost of holding inventory is greater than we observe with FIFO issuing. Consequently, the retailer maintains a lower service level on average (80% fill rate) without information sharing. Just as with FIFO issuing, when information is shared, it results in an improvement in the freshness of replenishment. Yet here, any improvement will not necessarily result in a decrease in product outdating because consumers buy the freshest product first. Moreover, any new replenishment may increase the likelihood of outdating product that is already in stock. Hence, with LIFO issuing the retailer is more constrained in its ability to take advantage of a fresher product replenishment. Thus, we find that the improvement in product freshness with information sharing is on average 11.2%, which is 20.5% less than that observed with FIFO issuing.

As a final point, it is interesting to note that retailer profitability generally increases by a much larger percentage by switching from LIFO to FIFO than from information sharing. Only in approximately 10% of the cases do we find that VOI is greater than switching from LIFO to FIFO and these instances, not surprisingly, correspond to where the VOI is greatest – low expected product lifetime, high variability in the age of receipts, high product cost, large batch size, and high demand variability. This result indicates that retailers who have not implemented FIFO may be better off trying to do so first, before making investments in information sharing.

5.2.2 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 in demand C_D, and the batch size Q. As shown in Table 5.2, the VOI increases with respect to both parameters. This occurs because the potential for product outdating increases when either demand is more uncertain or the replenishment batch size is larger. To demonstrate, consider an extreme case where $C_A = 0$ and $Q = 1$. Demand is deterministic and the retailer orders D units every period. In this scenario, the retailer fully satisfies demand and incurs no inventory related costs. Hence, VOI is zero.

Another related factor that impacts the retailer's ability to match supply and demand is the inherent uncertainty with respect to supply, as measured by the coefficient of variation in the age of replenishment C_A. We observe the same, if not stronger, relationship between supply uncertainty and the VOI as we do between demand uncertainty and the VOI. That is, the more uncertainty there is with regard to the age of replenishment the higher the VOI. Again, an extreme example is sufficient to demonstrate. Consider the case where $C_A = 0$. Here, there is no variability over time with respect to freshness of replenishment and hence the VOI is zero.

5.2.3 Drivers of Retail Profitability

Since $\pi^{NIS}(\boldsymbol{i})$ resides in the denominator of the VOI equation, we find that when $\pi^{NIS}(\boldsymbol{i})$ is small, the VOI is larger than when $\pi^{NIS}(\boldsymbol{i})$ is large. Clearly on a relative basis, even small changes will be more pronounced when the denominator is small. Hence, we find tat the VOI increases with respect to parameter values that induce a low level of profitability in the no information sharing case. Most notable among these include a low mean demand rate and a high product cost. For example, as the product cost increases from \$0.40 per unit to \$0.70 per unit, the average expected profit decreases from \$1.60 to \$0.58. This significant reduction in the denominator of the VOI equation is largely responsible for the increase in the VOI as the product cost increases. Notice that while the total improvement in expected profit due to information sharing also increases as the product cost increases (from \$0.06 to \$0.08), the change in the denominator clearly pronounces the relationship.

5.3 Impact on the Supplier

Our analysis would not be complete without an analysis on the impact of information sharing on the supplier. While the supplier is exogenous to the model, we can nevertheless measure the impact that information sharing has on its performance by considering the net change in expected retail orders. Across experiments, we observe a range of between –25.8% and +34.0% and a mean of –2.0% in the change in expected size of retail orders per period. The size of the change depends largely on the relative improvements in retailer outdating and retailer service. Improvements in retailer outdating translate to a decrease in orders to the supplier while improvements in service translate to an increase in orders to the supplier. In 27% of the experiments, the average expected order size to the supplier increases.

Generally, however, the reduction in unit outdating is greater than the increase in units of satisfied demand so that the supplier is worse off on average. In Table 5.4 we report the impact that information sharing has on the supply chain by reporting, at given percentiles across the 1,296 experiments, the change in orders to the supplier and the change in supplier revenue. We also report the impact on both the supplier's expected profit and the combined expected profit for both the retailer and the supplier. We do so by evaluating two cases: one in which the supplier's product margin is 10% and another in which it is 50%. This provides a relative comparison between cases when the supplier is a distributor (low margin) and a manufacturer (high margin).

Percentile	Supplier				Combined	
	Δ Order	Δ Revenue	Δ Profit (10% Margin)	Δ Profit (50% Margin)	Δ Profit (10% Margin)	Δ Profit (50% Margin)
0.00	-0.82	-0.54	-0.05	-0.20	0.00	-0.15
0.05	-0.36	-0.22	-0.02	-0.11	0.02	-0.01
0.10	-0.27	-0.14	-0.01	-0.07	0.02	0.00
0.25	-0.15	-0.07	-0.01	-0.04	0.03	0.02
0.50	-0.06	-0.03	0.00	-0.02	0.06	0.05
0.75	0.01	0.00	0.00	0.00	0.10	0.08
0.90	0.10	0.06	0.01	0.03	0.13	0.12
0.95	0.16	0.09	0.01	0.05	0.15	0.15
1.00	0.86	0.48	0.05	0.24	0.23	0.34

Table 5.4: Impact of information sharing on the supply chain

As sown in Table 5.4, while the supplier is worse off in a preponderance of the cases, total supply chain profits are *almost* always higher under information sharing. Only with the assumption of a high supplier product margin (and then only in 70 of the 1,296 experiments) do we find that the combined change in profitability of the retailer and supplier is negative. There is no clear pattern to these cases except that they all occur with a product cost of $w = \$0.70$, an expected lifetime of $\mu_A = 2$, and LIFO issuing. With a small retailer product margin and a large supplier product margin, a significant reduction in outdating can have a greater negative impact on the supplier than a positive impact on the retailer. Nevertheless, the viability of information sharing rests with the aggregate performance of hundreds if not thousands of perishable products and the results make clear that even with a high supplier product margin, total supply chain profitability increases in approximately 95% of the cases.

Since the supplier is generally worse off, some form of contract beyond the normal price-only contract is needed to induce the supplier to participate in the information sharing. Cachon and Lariviere (2003) discuss revenue sharing contracts, where the product is sold to the retailer at the supplier's cost and the retailer shares a pre-determined percentage of the revenue with the supplier. The percentage of revenue shared is generally set such that the contract is pareto improving. Given the widespread use of scanners in the retail industry, the cost of implementing and monitoring such a contract should not be prohibitive.

6. Heuristic Policies

In this section, we introduce, test, and explore the relative performance of heuristic policies. Clearly, the policies introduced in §3 are impractical to implement for larger sized problems given that the size of the state space expands exponentially with the age dependent vector of inventory. Hence, our purpose here is to develop policies that enable a broader evaluation on the VOI and that are more relevant to practice.

In §6.1 we define two heuristic policies that correspond to the two optimal policies for no information sharing and information sharing. In §6.2 we demonstrate through a series of tests that not only do these policies perform extremely well, they also qualitatively exhibit the same relative performance in terms of the VOI. Finally, in §6.3 we further our analysis on the VOI by using these policies in an extended numerical analysis of problems sized much larger than those explored in §5.

6.1 Heuristic Policies

The structure of the heuristic policies is predicated on a balance between simplicity and performance. Since a retailer can place an order each day and the lead-time is one day, the heuristics represent a myopic policy. That is, the order decision rests on whether sufficient stock exists in the current period that will carry over and minimize expected cost in the next period only. If sufficient stock exists, then the decision to order is postponed to the next day.

To begin, a major component of the heuristic policies is the one period loss function $L(I)$. This function does not take into account the cost of product outdating however. To do so, we introduce a new function. Let $outdate^t(\boldsymbol{i})$ denote the total estimated outdating cost associated with inventory $\boldsymbol{i}$, t days in the future from the current period, where $t \in \{1, 2, ..., M\}$. Hence, $outdate^1(\boldsymbol{i})$ denotes the total estimated future outdating cost of inventory $\boldsymbol{i}$ in the next period. Formally, we have

$$outdate^t(\boldsymbol{i}) = \begin{cases} w\sum_{D=0}^{i_1}(i_1 - D)\phi(D) + \sum_{D=0}^{I} outdate^{t+1}(\boldsymbol{i}')\phi(D) & 1 \le t \le M-1 \\ w\sum_{D=0}^{i_1}(i_1 - D)\phi(D) & t = M \end{cases}.$$

where

$$i'_x = \begin{cases} \left(i_{x+1} - \left(D - \sum_{z=1}^{x} i_z\right)^+\right)^+ & \text{FIFO issuing} \\ \left(i_{x+1} - \left(D - \sum_{z=x+2}^{M} i_z\right)^+\right)^+ & \text{LIFO issuing} \end{cases}.$$

With $outdate^t(\boldsymbol{i})$ and the one period loss function $L(I)$, we have the two essential components of our heuristic cost-to-go function $g(\boldsymbol{i})$ where

$$g(\boldsymbol{i}) = L(I) + outdate^1(\boldsymbol{i}).$$

Now, let $f^{HNIS}(\boldsymbol{i})$ denote the minimum total estimated cost for the Heuristic No Information Sharing policy (HNIS), where

$$f^{HNIS}(\boldsymbol{i}) = \min_{n \in N} \sum_{D=0}^{\infty}\sum_{A=1}^{M} g\left(\tau(\boldsymbol{i}, D, nQ, A)\right)\psi(A)\phi(D).$$

For the Heuristic Information Sharing policy (HIS), we simply add the age of replenishment to the state space so that

$$f^{HIS}(\boldsymbol{i}, A) = \min_{n \in N} \sum_{D=0}^{\infty} g\left(\tau(\boldsymbol{i}, D, nQ, A)\right)\phi(D).$$

While the heuristic policies, as defined, operate with the same state space as the optimal policies, they have two distinct advantages. First, the computational time is a fraction of that required for the corresponding optimal solution allowing larger problems to be solved. Second, it is not necessary to solve for all states simultaneously, except when evaluating the total expected cost in steady-state (i.e. performance evaluation). Hence, from a practical perspective, the policies are extremely fast and straight-forward to implement. Perhaps most importantly, they provide near optimal performance as we discuss below.

6.2 Comparison of Heuristic Policies to Optimal Policies

We test the heuristics by comparing their cost performance with the optimal policies. To do so, we use the full set of 1,296 numerical examples defined for the experimental design in §5 and employ the same solution procedure, except of course, we substitute the optimal policies with the heuristic policies for each state. We measure the performance of each heuristic policy by taking the percentage difference in expected profit, relative to the corresponding optimal policy. Overall, the results are very good. The no information sharing heuristic achieves, on average, a total expected profit that is 0.8% less than optimal and the information sharing heuristic achieves, on average, a total expected profit that is 1.7% less than optimal. We report the performance at selected percentiles of the 1,296 test cases in Table 6.1.

Percentile	No Info Heuristic	Info Sharing Heuristic
0.00	0.0%	0.0%
0.05	0.0%	0.2%
0.10	0.0%	0.4%
0.25	0.2%	0.9%
0.50	0.4%	1.5%
0.75	1.0%	2.3%
0.90	1.9%	3.4%
0.95	3.0%	4.3%
0.99	5.0%	6.5%
1.00	9.2%	9.8%

Table 6.1: Heuristic Performance

As shown in Table 6.1, the worst-case performance is less than 10% from optimality for both heuristics and is less than 5% from optimality in over 95% of the test cases. We were not able to identify any patterns in the results to explain why the performance under a few sets of parameter values was worse than others. In a second test, we compared the VOI achieved with the heuristics to that of the optimal policies. The average VOI of the heuristics across all 1,296 examples is 8.7%, or 0.9% less than reported for the optimal policies. This is not unexpected as the performance of the information sharing heuristic is also, on average, 0.9% further from optimality than the performance of the no information sharing heuristic. Hence, we would expect that the VOI to be underestimated by the heuristics. Moreover, a thorough comparison of the heuristic VOI to the optimal VOI, across parameter dimensions, demonstrates the same qualitative relationships. To demonstrate, in Table 6.2 we provide a side-by-side comparison of the average heuristic VOI and the average optimal VOI for all parameters and their values used in the test cases. From the basis of these comparisons, we consider the heuristic policies to be well suited for our purposes and proceed to support and extend our analysis on the VOI.

Parameter	Value	Heuristic VOI	Optimal VOI
Issue Policy	FIFO	9.6%	10.4%
	LIFO	7.7%	8.9%
Mean Demand	3	9.9%	10.8%
	4	7.4%	8.5%
Demand CV	.60	6.5%	7.5%
	.75	8.8%	9.8%
	.90	10.6%	11.6%
Expected Lifetime	2	10.6%	11.8%
	3	9.4%	9.8%
	4	6.0%	7.4%
Age CV	.2	3.2%	3.7%
	.3	7.7%	9.0%
	.4	15.0%	16.2%
Product Cost	0.40	2.9%	3.9%
	0.55	6.2%	7.4%
	0.70	16.8%	17.7%
Batch Size	1	6.5%	7.6%
	2	6.8%	7.6%
	4	7.3%	8.7%
	6	14.1%	14.7%

Table 6.2: Heuristic VOI compared to optimal VOI

6.3 Extended Analysis

We use the heuristics to explore the VOI in the context of problems sized much larger than those evaluated in §5. With the optimal policies and the bounds developed in §4, we are able to solve problems up to approximately four million states expediently. While this may sound considerable in itself, we are nevertheless limited to evaluating examples with small mean demand rates and small batch sizes. Here we explore problems with larger mean demand rates and a larger batch size. Specifically, we consider a set of experiments that comprise a factorial design for all combinations of the following parameters:

$$\mu_D \in \{5,6,7\} \qquad C_D \in \{0.5,0.6,0.7\} \qquad Q \in \{2,4,8\}$$

$$\mu_A \in \{2,3,4\} \qquad C_A \in \{0.2,0.3,0.4\} \qquad w \in \{0.4,0.55,0.70\}$$

As in the prior analysis, we duplicate the factorial design for each of the two issuing policies so that there are a total of 1,458 experiments. The age distribution $\psi(\cdot)$ in each experiment is determined randomly to achieve the chosen mean and coefficient of variation using the same procedure as described in §5. The maximum product lifetime M is 5 days and the holding cost rate is fixed at 25% of the product cost.

As one can see for the set of parameter values, we remain restricted in our choice of values due to the size of the state space because it is necessary to compute the heuristic policies for every viable state and then solve for steady state probabilities. To a large extent, computer memory, not computing speed is the limiting factor. Consider that with a product lifetime of five days, a mean demand rate of seven, and a coefficient of variation in demand equal to 0.7, there are approximately 32 million states to evaluate. Smart computing, using the bounds on the decision space as defined in §4 helps considerably and we can solve this problem in less than 20 minutes. Yet, it remains infeasible to extend our analysis to product lifetimes of six days. Consider that by doing so, the state space would exceed 888 million. This is not to say, however, that the heuristic policies could not be extended to these problems. In practice, it is only necessary to solve for a single state.

In general, the results on the VOI for this set of experiments provide the same insights as those in §5. The mean VOI across all 1,458 experiments is 4.1%, with a range between 0.0% and 64.0%. While the mean and range are smaller than observed with the optimal policies, this is not unexpected given both the larger mean demand rates we explore here and the performance of the heuristics themselves. In Table 6.3 we present the results of a sensitivity analysis where we report the average expected VOI for each parameter and value.

Parameter	Value	VOI
Issue Policy	FIFO	4.8%
	LIFO	3.5%
Mean Demand	5	4.5%
	6	3.8%
	7	3.7%
Demand CV	.50	3.4%
	.60	4.1%
	.70	5.4%
Expected Lifetime	2	6.4%
	3	3.6%
	4	2.4%
Age CV	.2	1.8%
	.3	3.8%
	.4	6.9%
Product Cost	0.40	1.5%
	0.55	3.3%
	0.70	7.6%
Batch Size	2	4.1%
	4	3.6%
	8	4.8%

Table 6.3: Sensitivity Analysis

Except for the smaller reported values for the VOI, qualitatively, the results are nearly identical to that observed in our sensitivity analysis of the optimal policies. The only difference we observe is that the VOI in this test, does not increase monotonically with respect to the batch size, although the VOI is largest for the largest batch size.

7. Conclusion

In this paper, we study the benefits of information sharing to a retailer that sells a perishable product with a fixed lifetime and is constrained to order in fixed lot sizes. We first propose a policy for the retailer under no information sharing and then provide an exact analysis for expected cost. We compare these results to ones obtained with information sharing using a numerical study and find that the retailer benefits the most from information sharing when: 1) the variability of the retailer's replenishment lifetime is high, 2) product lifetimes are short and batch order quantities are large, and 3) uncertainty in demand is significant and batch order quantities are large. We also find that information sharing is generally more beneficial when demand is satisfied with a FIFO issuing policy than with a LIFO issuing policy. In fact, we observe that it is generally more profitable to switch from LIFO to FIFO issuing (if possible) than from sharing information.

The average improvement from information sharing is 9.6%. The benefits of information sharing, however, are not necessarily shared with the supplier. Although the supplier is exogenous to our model, we observe that information sharing can result in a net decrease in retailer replenishment orders due to a reduction in the amount of retailer outdating. The benefits of information sharing to the whole supply chain however are almost always positive, indicating the possibility exist for pareto improvement through some form of coordination contract.

There are a number of important issues left to be addressed. First, we have not addressed the impact that holding fresher product has on retail sales. As mentioned earlier, the availability of perishable products is an order winning criteria of consumers. It is reasonable to assume that a firm carrying fresher product will observe higher demand and or higher margins than one that carries older product. We expect that the inclusion of this important relationship will only increase the value of information sharing. Other areas left for future research include models that address longer or stochastic lead times, random issuing policies, and non-stationary demand.

References

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.

Chiu, H.N. (1995), "An approximation to the continuous review inventory model with perishable items and lead times," *European Journal of Operational Research*, 28(2), 365-374.

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

Cachon, G. and M. Lariviere (2003), "Supply Chain Coordination with Revenue Sharing Contracts: Strengths and Limitations," Wharton School of Business Working Paper.

Chen, F. (2002), "Information Sharing and Supply Chain Coordination," To appear in the Handbook of Operations Research and Management Science: Supply Chain Management, by T. de Kok and S. Graves (eds.), North-Holland.

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

Chen, F. and B. Yu. (2001), Quantifying the Value of Leadtime Information in a Single-Location Inventory System. Working Paper, Columbia Business School.

Fries, B. (1975), "Optimal Ordering Policies for Perishable Inventory–II," *Operations Research*, Vol. 23, No. 1, 1975, 46-61.

Heller, T. (2002), "Sales Grow so does competition," *Progressive Grocer*, 81(10), 103-106.

Hennessy, T. (2000), "Growing Resentment," *Progressive Grocer*, 79(11), 69-74.

Huang, G., J. Lau, and K. Mak (2003), "The impacts of sharing production information on supply chain dynamics: a review of the literature," *International Journal of Production Research*, ***41***(7), 1483-1518.

Ketzenberg, M. and M. Ferguson (2003), Sharing Information to Manage Perishables, College of Business, Colorado State University working paper.

Kollars, D (2001), "A changing landscape," *Progressive Grocer*, 80(9), 49-54.

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

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

Leiberman (1958), "LIFO vs. FIFO in Inventory Depletion Management," *Management Science*, Vol. 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*, Vol. 23, No. 4, 1975, 735-749.

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.

Nandakumar, P. and T. Morton (1993), Near Myopic Heuristic for the Fixed-Life Perishability Problem, *Management Science,* 39(12), 1490–1498.

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

Pfankuch, B. (2004), Store Manager of Albertson's, Fort Collins, CO, Personal Interview.

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

Raper, B. (2003), Divisional Distribution Manager of Mayfield Dairy, Personal Interview.

Sahin, F. and E. P. Robinson (2002), "Flow Coordination and Information Sharing in Supply Chains: Review, Implications, and Directions for Future Research," *Decision Sciences*, 33(4), 1-32.

Souza, G., M. Ketzenberg, and D. Guide, "Capacitated Remanufacturing with Service Level Constraints," *Production and Operations Management*, Vol. 11, Summer 2002, 231-248.

Zipkin, P. (2000), Foundations of Inventory Management, Irwin