Appendix to “Managing Slow Moving Perishables in the Grocery Industry”

Michael Ketzenberg
Department of Information & Operations Management
Mays Business School
322 Wehner Building
4217 TAMU
Texas A&M University
College Station, Texas 77843-4217

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

May, 2007

A1: Retailer Order Probabilities in the NIS Case

Here, we characterize the distribution $\psi_{\tilde{D}}(\beta)$ 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 et al. (2007) 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, …, 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,

$$\tilde{D} = N(Q) + 1 \equiv \min\{n : S_n = X_1 + X_2 + \cdots + X_n > Q\}, \qquad \text{(A1)}$$

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. Thus, we let $\psi_{\tilde{D}}(\beta)$ represent the cdf of $\tilde{D}$ with a mean of Q/μ and a variance of $Q\sigma^2/\mu^3$.

A2: Solution Procedure for the DIS Case

```
PROCEDURE f(\vec{i}, A)
  FOR q = 0 TO Q STEP Q                         ;Evaluate q = 0 (1st)and q = Q (2nd).
     Profit:= G(I) - q(1 - m_0)                 ;Initialize profit to one period profit.
     IF (q>0) or (A=0)  THEN                    ;If supplier has no inventory going
        Determine λ                             ;  into next period, determine λ.
     ELSE                                       ;if supplier has inventory going into
        λ:= 0                                   ;next period, then no supplier order.
     FOR D = 0 TO MAX DEMAND                    ;Evaluate all realizations of demand.
        Profit = Profit + f(τ(\vec{i}, D, q, A), A')φ(D)  ;add in future expected profit.
     ENDFOR (D)
     IF q<Q THEN                                ;if 1st time through, then save results
        BEGIN                                   ;for later comparison to q = Q.
          SaveProfit:=Profit
          SaveLambda:= λ
        END
     ELSE                                       ;2nd time through, compare profit
        IF Profit< SaveProfit THEN              ;of q=0 (Saveprofit) to q = Q (Profit).
          BEGIN                                 ;Case q = 0 > q = Q.
             q*:=0                              ;Set optimal decisions and
             f(\vec{i}, A):= SaveProfit         ;expected profit.
             λ*:= SaveLambda
          END
        ELSE                                    ;Case q = Q > q = 0.
          BEGIN
             q*:=0                              ;Set optimal decisions and
             f(\vec{i}, A):=Profit              ;expected profit.
             λ*:= λ
          END
  ENDFOR (q)
ENDPROCEDURE
```

A3: Sensitivity Analysis

Generally, we find that the VOI and the VCC 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 sensitivity to each parameter in Figure 1. The height of each bar corresponds to the average VOI and 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 section A6.

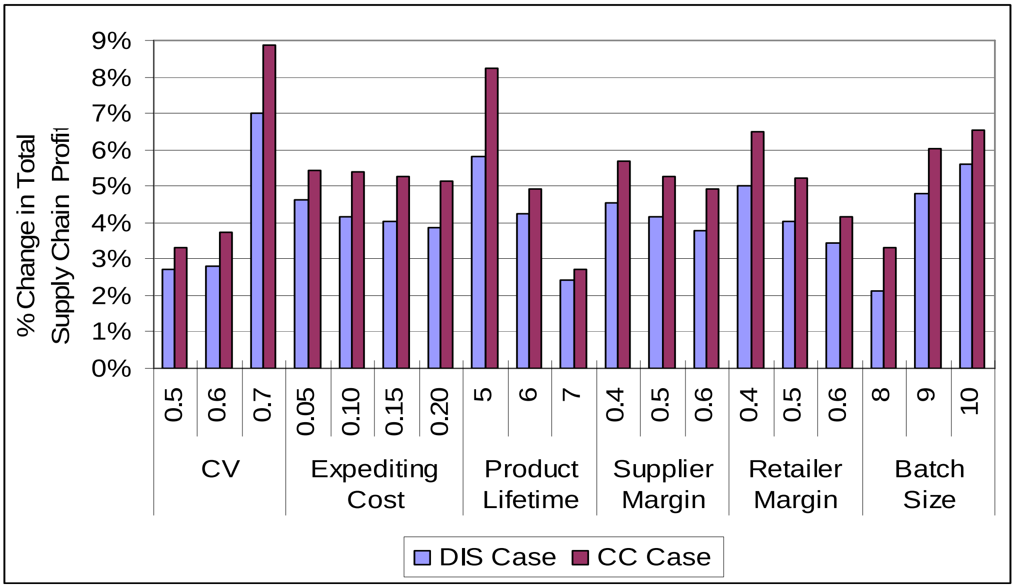

Figure 1: Sensitivity of the average VOI/VCC for each fixed parameter value

Product Perishability

As shown in Figure 1, the VOI and the VCC both 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, resulting 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. As an example, 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 2, we show through an illustrative example 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 $m_0 = 0.6$.

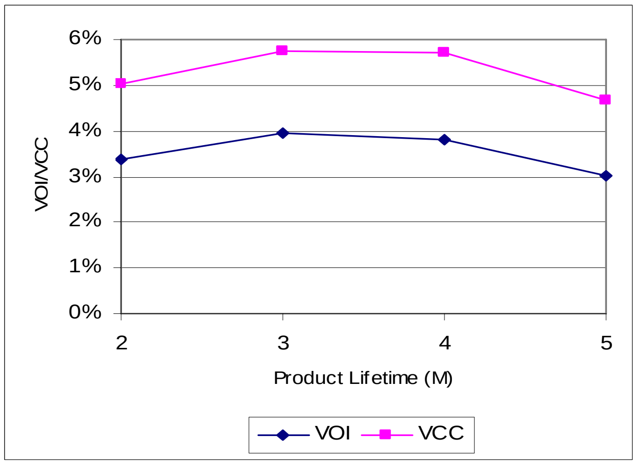

Figure 2: Sensitivity of the average VOI/VCC for product lifetime

Long product lifetimes result in small VOI and VCC because the prospect of outdating is small. In this scenario, service levels are higher and outdating is lower so any improvement in product freshness does 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, both the VOI and the VCC were trivial: the average is 0.1% and the maximum is 1.3%.

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 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 1 where the values for the VOI and VCC are averaged across experiments at each level of *M* and *Q*. It is quite clear both the VOI and VCC quickly approach zero as the batch size approaches the mean demand rate.

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

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

Table 1: Average VOI (left) and VCC (right) with respect to small batch sizes (*Q*)

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, the VOI and the VCC increase. We show these relationships in Table 2 where the values for the VOI and the 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 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.

A4: Extensions

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

Price Sensitivity to Freshness and Minimum Product Freshness

In our earlier analysis, we assume that supplier receives the same revenue per unit, regardless of its product freshness, and 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

$$A_{\min} = \min\left(A : \sum_{D=0}^{\infty}\left[-p_{1,A}(Q-D)^+ - h_0 A\left(\frac{Q-(Q-D)^+}{2}\right) + \left(p_0 - p_{1,A}\right)\min(Q,D)\right]\phi_A(D) > 0 \right). \quad \text{(A2)}$$

On the right side of (A2), *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 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 $P(\tilde{D} \ge M - A_{\min} + \beta)$. 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 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 section A5. 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) 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 3 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 3: VOI and VCC at percentiles for each value of δ

As shown in Table 3, 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, 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 for $Q \in (10,11)$. 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 4 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 4: % Share of value in the DIS Case for each value of δ

Note in Table 4 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.

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, it can be used to find the optimal Q through a full enumeration search 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 the previous section 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 3, 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 3 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.

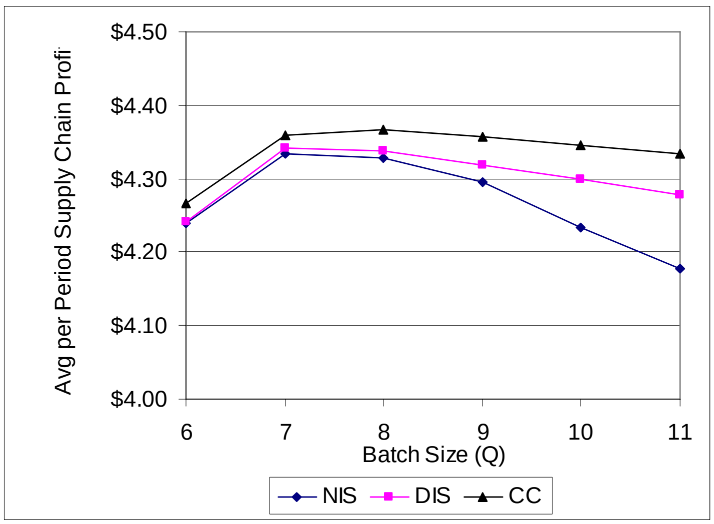

Figure 3: Average Total Profit at each value of Q

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, based on our limited numerical results, it appears that information sharing and centralized control are less valuable if a supply chain can choose the optimal batch size.

A5: Supplier's Policies with Model Extensions

In the previous section, our extension to freshness dependent pricing for the supplier fundamentally changes 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 $\Omega_{\tilde{D}}(\beta)$ 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:

$$\max_{\alpha}\left(\sum_{\beta=1}^{M}\begin{cases}\left[Q\left(p_{1,M}-c_1\right)-b\right]\Omega_{\tilde{D}}(\beta) & \alpha \geq \beta \\ Q\left[\left(p_{1,M-\beta-\alpha+1}-c_1\right)-h_1\left(\beta-\alpha-1\right)\right]\Omega_{\tilde{D}}(\beta) & \alpha < \beta, M-\beta+\alpha+1 \geq A_{\min} \\ Q\left[\left(p_{1,M-\beta+\alpha-A_{\min}+2}-2c_1\right)-h_1\left(\beta-\alpha-1\right)\right]\Omega_{\tilde{D}}(\beta) & \alpha < \beta, M-\beta+\alpha+1 < A_{\min}\end{cases}\right). \text{(A3)}$$

The expectation of the suppliers profit (A3) 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 $\Omega_{\tilde{D}}(\beta)$. 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 $\Omega_{\tilde{D}}(\beta)$ in the same manner as the distribution $\psi_{\tilde{D}}(\beta)$ expressed in section A1. 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. Resolving over multiple iterations may still 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 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.

A6: 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.

References

Bai, L., C. Alexopoulos, M. Ferguson, K. Tsui. 2007. A simple and robust batch-ordering inventory policy for unobserveable demand. Working paper, H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta.

Feller, W. 1949. Fluctuation theory of recurrent events. *Trans. Amer. Math. Soc.* 67 98–119.

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

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