11-1999

Competitive and Cooperative Inventory Management in a Two-Echelon Supply Chain With Lost Sales

Gerard. P. Cachon
University of Pennsylvania

Competitive and Cooperative Inventory Management in a Two-Echelon Supply Chain With Lost Sales

Abstract

This paper studies inventory management in a two echelon supply chain with stochastic demand and lost sales. The optimal policy is evaluated and compared with the competitive solution, the outcome of a game between a supplier and a retailer in which each firm attempts to maximize its own profit. It is shown that supply chain profit in the competitive solution is always less than the optimal profit. However, the magnitude of the competition penalty is sometimes a trifle, sometimes enormous. Several contracts are considered to align the firms' incentives so that they choose supply chain optimal actions. These contracts contain one or more of the following elements: a retailer holding cost subsidy (which acts like a buy-back/return policy), a lost sales transfer payment (which acts like a revenue sharing contract) and inventory holding cost sharing. With the latter each firm incurs a fixed fraction of the total supply chain holding cost. It is found that the retailer holding cost subsidy is generally not sufficient to coordinate the supply chain. The most effective contract combines a lost sales transfer payment with inventory holding cost sharing: it always coordinates the supply chain, both players are always better o¤ and it is simple to evaluate.

Disciplines

Business Administration, Management, and Operations | Operations and Supply Chain Management

Competitive and Cooperative Inventory Management in a Two-Echelon Supply Chain with Lost Sales*

Gérard P. Cachon
The Fuqua School of Business · Duke University · Durham · NC · 27708
gpc@mail.duke.edu · www.duke.edu/~gpc

November 1999

Abstract

This paper studies inventory management in a two echelon supply chain with stochastic demand and lost sales. The optimal policy is evaluated and compared with the competitive solution, the outcome of a game between a supplier and a retailer in which each firm attempts to maximize its own profit. It is shown that supply chain profit in the competitive solution is always less than the optimal profit. However, the magnitude of the competition penalty is sometimes a trifle, sometimes enormous. Several contracts are considered to align the firms' incentives so that they choose supply chain optimal actions. These contracts contain one or more of the following elements: a retailer holding cost subsidy (which acts like a buy-back/return policy), a lost sales transfer payment (which acts like a revenue sharing contract) and inventory holding cost sharing. With the latter each firm incurs a fixed fraction of the total supply chain holding cost. It is found that the retailer holding cost subsidy is generally not sufficient to coordinate the supply chain. The most effective contract combines a lost sales transfer payment with inventory holding cost sharing: it always coordinates the supply chain, both players are always better off and it is simple to evaluate.

*Thanks is extended to Martin Lariviere for his many helpful comments. Since I will soon present this work at the Kenan-Flagler School, University of North Carolina and the Kellogg School of Management, Northwestern University, I thank those seminar participants in advance.

Essentially all supply chains operate as a collection of independent agents, each responsible for managing a subset of the supply chain, each with its own objectives. While it is unlikely that the agents' objectives are entirely orthogonal, one should not expect harmoniously aligned incentives either. But, when do those difference lead to sub-optimal supply chain performance? As in a good marriage, some disagreement is natural, relatively harmless and certainly does not justify costly intervention/counseling. Only when the supply chain operates significantly below its optimal performance should the firms seek therapy, i.e., align their incentives. But what kinds of agreements align the firms' incentives so that (1) each firm can be reasonably certain that the other firm will choose a supply chain optimal action (i.e., uphold its part of the deal) and (2) each firm is better off with the agreement than without it, i.e., it is Pareto improving? This paper studies those questions with a two firm supply chain inventory model.

The two firms, a supplier and a retailer, operate over an infinite horizon. The supplier produces a single product at a finite production rate and at a fixed marginal production cost. The retailer purchases the product at a fixed wholesale price and sells the product at a fixed retail price. Both the supplier and the retailer incur holding costs on their inventory. The inter-arrival times of customers are stochastic and each customer purchases exactly one unit if a unit is available at the retailer. If a unit is not available, that potential sale is lost.

Each firm attempts to maximize its expected profit per unit time knowing that the other firm has the same objective. The firms could choose actions that lead to optimal supply chain performance: the optimal policy belongs to their feasible set of actions and all information is common knowledge (so the firms are able to evaluate the optimal policy). Nevertheless, the firms may choose sub-optimal policies because their incentives are not aligned with the supply chain's incentives.

It is not obvious whether the firms will carry too much or too little inventory relative to the optimal amount. Increasing supply chain inventory benefits both firms, so the supply chain enjoys a larger benefit from increasing inventory than either firm, which suggests the firms will not stock enough. But neither firm incurs the full cost of increasing supply chain inventory (the increase may be divided between the firms), which suggests the firms will stock too much. If those contradictory incentives balance each other, the firms may in fact choose policies that are reasonably close to optimal, if not optimal. If not, it is important to identify which incentive dominates so that a contract can be constructed that aligns the firms' incentives: a contract that must increase inventory is obviously different than a contract that must reduce inventory.

The next section reviews the related literature. §2 details the model. §3 provides the supply chain optimal solution. §4 analyzes the competitive supply chain inventory game. §5 presents a numerical study that investigates the magnitude of the competition penalty (the decrease in supply chain performance due to decentralized decision making). §6 investigates how the firms can better align their incentives. Several contracts are considered. §7 summarizes and discusses the results.

1. Literature Review

This work is similar to Cachon and Zipkin (1999), hereafter referred to as CZ. They study a two echelon serial supply chain with stochastic consumer demand and complete backordering, i.e., there are no lost sales. Both firms incur holding costs and a backorder penalty per unit of time for each unit that is backordered at the retailer; the supplier is not charged for its own backorders, it is only charged when units are backordered at the retailer. That fee reflects the supplier's desire to maintain an adequate stock of its product at the retailer. This model is also a two echelon serial supply chain with stochastic consumer demand. But when a customer arrives at the retailer and the retailer has no stock, a lost sale occurs. Hence, as in CZ, both firms are concerned about the availability of inventory at the retailer, but in this model stock outs create opportunity costs (each firm's margin on the product) rather than backorder penalties. A second significant difference is that there is no replenishment delay between the two echelons in this model, whereas there is a fixed replenishment delay between the two echelons in CZ. (In both models the upper echelon experiences a replenishment delay.) The implication of this difference is discussed later.

This work also resembles Caldentey and Wein (1999), hereafter referred to as CW. They study a two firm serial supply chain in which the supplier controls the production rate of his manufacturing facility and the retailer chooses an inventory stocking policy. Both the production and demand inter-arrival times are exponentially distributed. All demands are backordered, and, as in CZ, both firms incur a backorder penalty cost for backorders at the retail level. In this model production inter-arrival times are also exponentially distributed, but the rate is exogenous. Hence, as in CZ, both the supplier and the retailer manage inventory, whereas only the retailer manages inventory in CW. As in this model, CW assume units are transferred immediately between the supplier and the retailer once production is completed.

Both CZ and CW consider linear transfer payment contracts to coordinate the supply chain. This work considers similar contracts.

There is other research that studies competitive supply chain inventory management over an infinite horizon. Cachon (1997) extends the model in CZ to include multiple retailers and batch ordering. Chen (1999), Lee and Whang (1999) and Porteus (1997) focus on contracts to coordinate serial supply chains with complete backordering. Chen, Federgruen and Zheng (1997) study competition and coordination in a multiple-retailer model with deterministic demand, so there are no lost sales.

There are numerous papers that consider supply chain contracting over short horizons (see Tayur, Ganeshan and Magazine, 1999). These papers are generally based on the newsvendor model, and so lost sales are possible. Pasternack (1985) is a classic example. He investigates a model with a single supplier and a single retailer and shows that the retailer purchases too little inventory with a simple wholesale price contract. The supplier can increase the retailer's order quantity with the use of a buy-back contract. (A buy-back contract specifies a price that the supplier will purchase left over stock from the retailer.) Not only can a buy-back contract induce the retailer to choose the supply chain optimal order quantity, it can also arbitrarily divide supply chain profits between the two firms. With the similar models, Lariviere and Porteus (1999) study the performance of a wholesale price only contract and Tsay (1999) studies the performance of a quantity flexibility contract.

Plambeck and Zenios (1999) study a model in which a principal offers a contract to an agent that controls a Markov decision process, hence only one player in their game directly controls the performance of the system. In this model system performance depends on the actions of both players, and either player may offer a contract.

There are many papers that investigate inventory competition among a group of retailers: e.g., Anupindi, Bassok and Zemel (1999a,b), Li (1992), Lippman and McCardle (1997), Mahajan and van Ryzin (1999), Wang and Gerchak (1999). Those models are different in that they include demand spillovers: if one firm is out of stock then the demand at the other firms increases. There are no demand spillovers in this model since there is only one retailer. Furthermore, those model assume a perfectly reliable inventory source for all retailers, whereas in this model the reliability of a retailer's source depends on the endogenous actions of the supplier.

2. The Model

This section describes a game between a supplier and a retailer. All information is common knowledge: each firm knows the rules of the game, their own costs, their opponent's costs, etc.

The supplier sells a single product to the retailer at a fixed wholesale price per unit, w. Customer inter-arrival times at the retailer are exponentially distributed with rate λ. Each customer purchases exactly one unit, if a unit is available at the retailer when the customer arrives. The purchase price is r per unit, $r - w \geq 0$. If the retailer is without inventory when a customer arrives, then the customer departs immediately and never returns; that potential sale is lost.

The supplier replenishes its inventory with a production process that has exponentially distributed inter-production times with rate μ. Define $\rho = \lambda/\mu$ and assume $\rho < 1$. (Allowing $\rho > 1$, while possible, generates few additional qualitative insights.) Once production of a unit is completed, the unit is immediately part of the supplier's inventory. The supplier is the retailer's only source of inventory. There is no time delay to transfer a unit of inventory from the supplier to the retailer, i.e., the retailer receives any order instantly as long as the supplier has inventory to satisfy the order. Hence, the supplier's production process is the only replenishment delay in this supply chain.

There are some settings in which the single replenishment delay is not onerous. In the personal computer industry suppliers (PC assemblers) and retailers (distributors) often co-locate in the same facility, in which case the lead time between the two firms is just the time to move a box between two rooms. Alternatively, the zero replenishment lead time is a reasonable approximation when the shipping time between the supplier and the retailer is relatively short so that the retailer generally does not stock out when there are units in-transit, e.g., if shipping occurs over night but demand occurs only during the day. While appropriate in a broader set of circumstances, a replenishment delay between the supplier and the retailer would introduce significant analytical challenges: the supply chain optimal policy would depend on how inventory is allocated within the system, rather than on just the total amount of inventory in the system. Those challenges notwithstanding, introducing a replenishment delay is an important extension for future research.

The supplier incurs a production cost c per unit produced, $w - c \geq 0$, and a holding cost h per unit in inventory per unit time. There are no holding costs for units in the production process. The retailer also incurs a holding cost at rate h per unit. (If the retailer's holding cost were higher than the supplier's holding cost, then the optimal solution has the supplier holding all inventory. A model with different holding costs is only interesting if there is a replenishment delay between the supplier and the retailer.)

A lost sale generates a $(r - w)$ opportunity cost for the retailer and a $(w - c)$ opportunity cost for the supplier, but a lost sale might generate other, indirect, consequences for the firms. For example, the supplier might lose some good will with the customer or the customer might switch to another product the retailer carries. To model those consequences, each lost sale creates a b_s charge to the supplier and a b_r charge to the retailer. To avoid trivial situations, assume $b_s \geq -(w - c)$ and $b_r \geq -(r - w)$: either firm could experience an indirect benefit due to a lost sale (e.g., the retailer may sell another product to the customer), but that benefit would not compensate the firm for its opportunity cost of a lost sale. In other words, each firm has some *a priori* incentive to avoid lost sales. Let $b = b_s + b_r$.

To manage inventory each firm uses a base stock policy: when the supplier's inventory is less than S_s the supplier maintains production, otherwise production is idle; and the retailer orders a unit of inventory from the supplier whenever its inventory is less than S_r. Given that the two firms are using base stock policies, the supplier has inventory only when the retailer's inventory is S_r. The supplier produces when its inventory is less than S_s, i.e., when the supply chain's inventory is less than $S = S_s + S_r$, and ceases production when its inventory is S_s, i.e., when the supply chain's inventory is S. So a centralized controller of the supply chain generates the same replenishment decisions as the firms when the controller uses a base stock policy with base stock level S. In fact, a base stock policy is optimal for the supply chain (i.e., it maximizes supply chain profit per unit time). This claim is made without formal proof, since a brief verbal argument should suffice: if production is initiated when there are x units in stock, then the expected profit from that unit in production is decreasing in x; so production should cease for all x greater than or equal to some S.

Let S^o be the supply chain's optimal base stock level. So even with decentralized operations the firms can optimize the supply chain as long as they choose $S_s + S_r = S^o$. Furthermore, it can be shown that a base stock policy is optimal for each firm assuming the other firm is using a base stock policy. Hence, base stock policies are quite reasonable for this setting.

The firms simultaneously choose their base stock levels (i.e., they must choose their base stock level before observing the other firm's base stock level) and then they are committed to those base stock levels over an infinite horizon (i.e., this is a one-shot game). A firm's only choice variable is its base stock level, and each firm must choose a base stock level: exiting from the game is not an option. Each firm chooses its base stock level with the objective to maximize its expected profit per unit time. For technical reasons, and without loss of generality, assume the firms' base stock levels are chosen from an interval: for $i \in \{s, r\}$, $S_i \in [0, \widehat{S}]$, where $\widehat{S}$ is a very large constant. A firm's base stock level will also be referred to as its strategy. A pair of base stock levels, $\{S_s^*, S_r^*\}$, is a Nash equilibrium if neither firm has a profitable unilateral deviation, i.e., each firm chooses a best response to the other firm's strategy.

It remains to evaluate profit functions for the supply chain, the supplier and the retailer. Begin with the supply chain. With base stock level $S \in \{0, 1, ..., \widehat{S}\}$ the supply chain's inventory belongs to the set $\{0, 1, ..., S\}$. Say the supply chain is in state i when there are i units of inventory in the supply chain. The supply chain's inventory is a continuous time Markov chain with $S + 1$ states, and, more specifically, it is a birth and death process. Let p_i be the probability the supply chain has i units of inventory in steady state. It is not difficult to show that, for $S \in \{0, 1, 2, ...\}$,

$$p_i(S) = \frac{\lambda^{S-i}\mu^i}{\sum_{j=0}^{S} \lambda^{S-j}\mu^j} = \frac{1-\rho}{\rho^{i-S} - \rho^{i+1}}.$$

(Note that $p_i(S)$ is the steady state probability that there are i customers in a $M/M/1/K$ queue with arrival rate μ, demand rate λ, and $K = S$, Kleinrock, 1975.) Define $P_i(S) = \sum_{j=0}^{i} p_j(S)$, i.e., $P_i(S)$ is the probability there are i or fewer units in the supply chain with base stock policy S:

$$P_i(S) = \frac{\rho^{S-i} - \rho^{S+1}}{1 - \rho^{S+1}}, \quad S \in \{0, 1, 2, ...\}.$$

Let $D(S)$ be the supply chain's average sales rate (in units). So, $\lambda - D(S)$ is the lost sales rate. Let $I(S)$ be the supply chain's average inventory. It follows, again for $S \in \{0, 1, 2, ...\}$, that

$$D(S) = \lambda(1 - P_0) = \lambda \left(\frac{1 - \rho^S}{1 - \rho^{S+1}} \right), \tag{1}$$

and

$$I(S) = S - \sum_{j=0}^{S-1} P_j = \frac{S - \left(\frac{\rho}{1-\rho}\right)\left(1 - \rho^S\right)}{1 - \rho^{S+1}}. \tag{2}$$

The retailer's average inventory is

$$I_r(S_s, S_r) = S_r - \sum_{j=0}^{S_r - 1} P_j = \frac{S_r - \left(\frac{\rho}{1-\rho}\right)\left(\rho^{S_s} - \rho^S\right)}{1 - \rho^{S+1}} \tag{3}$$

and the supplier's average inventory is

$$I_s(S_s, S_r) = I(S) - I_r(S_s, S_r) = \frac{S_s - \left(\frac{\rho}{1-\rho}\right)\left(1 - \rho^{S_s}\right)}{1 - \rho^{S+1}} \tag{4}$$

Note that the right hand sides of (1), (3) and (4) are continuous and differentiable in S_s and S_r. Thus, while the analysis of this model could be done assuming S_s and S_r are restricted to integer values (and therefore S is also integer), it is analytically more convenient to assume that the above inventory and sales rate functions apply for all $S_s \in [0, \widehat{S}]$ and $S_r \in [0, \widehat{S}]$. (The actual sales and inventory rates for the supply chain are $D(\lceil S \rceil)$, $I(\lceil S \rceil)$, $I_s(\lceil S_s \rceil, \lceil S_r \rceil)$ and $I_r(\lceil S_s \rceil, \lceil S_r \rceil)$, where $\lceil x \rceil$ is the smallest integer that is greater than or equal to x.) The qualitative impact of this assumption is minimal: with the continuous approximation it will be shown that there is a unique optimal base stock level for the supply chain and each firm has a unique base stock level given the base stock level chosen by the other firm, but those uniqueness results are lost if the base stock levels are restricted to the set of non-negative integers.

Let $\pi(S)$, $\pi_s(S_s, S_r)$ and $\pi_r(S_s, S_r)$ be the supply chain's, the supplier's and the retailer's average profit per unit time respectively:

$$\begin{aligned}
\pi(S) &= (r - c)D(S) - hI(S) - (b_r + b_s)(\lambda - D(S)) \\
\pi_r(S_s, S_r) &= (r - w)D(S_s + S_r) - hI_r(S_s, S_r) - b_r(\lambda - D(S_s + S_r)) \\
\pi_s(S_s, S_r) &= (w - c)D(S_s + S_r) - hI_s(S_s, S_r) - b_s(\lambda - D(S_s + S_r))
\end{aligned}$$

Let $m_r = r - w + b_r$, $m_s = w - c + b_s$, and $m = m_r + m_s$. Those constants are referred to as margins, since they represent the difference in profit between making a sale and generating a lost sale. Given the bounds on the lost sales penalties, each margin is non-negative. The profit functions can now be written as

$$\begin{aligned}
\pi(S) &= mD(S) - hI(S) - (b_r + b_s)\,\lambda \\
\pi_r(S_s, S_r) &= m_r D(S_s + S_r) - hI_r(S_s, S_r) - b_r\lambda \\
\pi_s(S_s, S_r) &= m_s D(S_s + S_r) - hI(S_s + S_r) + hI_r(S_s, S_r) - b_s\lambda
\end{aligned}$$

3. The Centralized (Optimal) Supply Chain Solution

The section evaluates the centralized (optimal) supply chain solution.

Theorem 1 *$\pi(S)$ is strictly quasi-concave.*

Proof. Differentiate :

$$D'(S) = \lambda\frac{(1-\rho)\rho^S \ln(1/\rho)}{(1-\rho^{S+1})^2}, \quad I'(S) = \frac{1-\rho^{S+1} - (S+1)\rho^{S+1}\ln(1/\rho)}{(1-\rho^{S+1})^2}$$

and

$$\begin{aligned}
\pi'(S) &= mD'(S) - hI'(S) \\
&= \frac{h}{(1-\rho^{S+1})^2}\left(\left(\frac{m(1-\rho)\ln(1/\rho)\,\lambda}{h\rho} + 1\right)\rho^{S+1} - 1 + (S+1)\rho^{S+1}\ln(1/\rho)\right)
\end{aligned}$$

Both ρ^{S+1} and $(S+1)\rho^{S+1}$ are decreasing in S. So there is some $\widetilde{S}$ such that $\pi'(S) \leq 0$ for all $S \geq \widetilde{S}$ and $\pi'(S) > 0$ for all $S < \widetilde{S}$, i.e., $\pi(S)$ is strictly quasi-concave.□

Given Theorem 1, there is a unique optimal base stock level, S^o, which is the solution to the first order condition:

$$mD'(S^o) = hI'(S^o). \tag{5}$$

S^o is easy to evaluate numerically. Let $\pi^o = \pi(S^o)$.

The following lemma indicates that the supply chain optimal inventory is positive only if $m\lambda$ is sufficiently large relative to h, holding ρ constant. It is rather uninteresting to study a supply chain that cannot justify stocking at least some inventory, so assume (6) holds throughout.

Lemma 2 *$S^o > 0$ only if*

$$\frac{m\lambda}{h} > \frac{1}{\ln(1/\rho)} - \frac{\rho}{1-\rho} \tag{6}$$

Proof. Since $\pi(S)$ is strictly quasi-concave $S^o > 0$ only if $\pi'(0) > 0$, which yields the above condition.□

When $b > 0$ it is possible that $\pi(S^o) < 0$ even if $S^o > 0$. It is not clear why a firm would participate in that market, nevertheless, positive profit constraints are not imposed in this model. (Those constraints are also referred to as participation constraints.) CW do impose those constraints in their model.

4. Decentralized Inventory Game Analysis

This section assumes that the supplier and the retailer choose their own inventory policies with the objective of maximizing their own profit. The main question is whether the firms choose policies that maximize total supply chain profit?

The first step in the game analysis determines each firm's optimal strategy choice assuming the other firm's strategy choice is known and fixed. Then, the existence of Nash equilibria is demonstrated. Finally, the Nash equilibria and the supply chain optimal solution are compared.

Let

$$\begin{aligned} S_r^*(S_s) &= \arg\max_x \pi_r(S_s, x) \\ S_s^*(S_r) &= \arg\max_x \pi_s(x, S_r) \end{aligned}$$

Based on the following theorems, $S_r^*(S_s)$ and $S_s^*(S_r)$ are functions, i.e., each firm always has a unique best response to the other firm's strategy choice.

Theorem 3 $\pi_r(S_s, S_r)$ *is strictly quasi-concave in* S_r.

Proof. Differentiate,

$$\frac{\partial \pi_r(S_s, S_r)}{\partial S_r} = m_r \lambda \frac{(1-\rho)\rho^S \ln(1/\rho)}{(1-\rho^{S+1})^2} - h \frac{\left(1-\rho^{S+1}\right) - \rho^{S+1}\left(S_r + \frac{1-\rho^{S_s+1}}{1-\rho}\right)\ln(1/\rho)}{\left(1-\rho^{S+1}\right)^2}. \tag{7}$$

The sign of (7) is the same as the sign of the following expression:

$$\left(\frac{m_r \lambda}{h}\right)(1-\rho)\rho^S \ln(1/\rho) - \left(1-\rho^{S+1}\right) + \rho^{S+1}\left(S_r + \frac{1-\rho^{S_s+1}}{1-\rho}\right)\ln(1/\rho) \tag{8}$$

Hence $\pi_r(S_s, S_r)$ is strictly quasi-concave in S_r if there exists a $\widetilde{S}_r \geq 0$ such that (8) is positive for all $S_r < \widetilde{S}_r$ and non-positive for all $S_r \geq \widetilde{S}_r$. When $\rho < 1$, the derivative of (8) with respect to S_r is negative,

$$-\rho^{S+1}\ln(1/\rho)^2\left(\left(\frac{m_r \lambda}{h}\right)\left(\frac{1-\rho}{\rho}\right) + \frac{1-\rho^{S_s+1}}{1-\rho} + S_r\right),$$

and so the result is confirmed.□

Theorem 4 $\pi_s(S_s, S_r)$ *is strictly quasi-concave in* S_s.

Proof. Differentiate,

$$\begin{aligned}
\frac{\partial \pi_s(S_s,S_r)}{\partial S_s} &= m_s\lambda\frac{(1-\rho)\rho^S\ln(1/\rho)}{(1-\rho^{S+1})^2} - h\frac{1-\rho^{S+1}-(S+1)\rho^{S+1}\ln(1/\rho)}{(1-\rho^{S+1})^2} \\
&\quad +h\frac{\rho^{S_s+1}\left(1-(S_r+1)\rho^{S_r}+S_r\rho^{S_r+1}\right)\ln(1/\rho)}{(1-\rho)\left(1-\rho^{S+1}\right)^2} \\
&= h\frac{z\rho^{S+1}-1+(S+1)\rho^{S+1}\ln(1/\rho)}{\left(1-\rho^{S+1}\right)^2}
\end{aligned}$$

where

$$z=\left(\frac{m_s\lambda}{h}\right)\frac{(1-\rho)\ln(1/\rho)}{\rho}+1+\frac{\left(1-(S_r+1)\rho^{S_r}+S_r\rho^{S_r+1}\right)\ln(1/\rho)}{(1-\rho)}>0.$$

Since both ρ^{S+1} and $(S+1)\rho^{S+1}$ are decreasing in S_s, there exists an $\widetilde{S}_s$ such that $\pi_s(S_s,S_r)$ is increasing for all $S_s\leq\widetilde{S}_s$ and decreasing otherwise, i.e., $\pi_s(S_s,S_r)$ is strictly quasi-concave in S_s.□

While the next Theorem confirms the existence of at least one Nash equilibrium, in fact, it is possible that multiple Nash equilibria exist. §5 provides data on the frequency of multiple Nash equilibria.

Theorem 5 *A Nash equilibrium $\{S_s^*,S_r^*\}$ exists in the decentralized inventory game.*

Proof. From Theorem 2.4 in Friedman (1986), a Nash equilibrium exists if: (1) the players have compact and convex strategies; (2) each player's payoff function is defined, continuous and bounded for all possible strategies; and (3) each player's payoff function is unimodal in its strategy. It is straightforward to confirm the continuity of each player's payoff function. The payoff functions are bounded because there exists a maximum supply chain profit. Theorems 3 and 4 confirm the third condition. If the continuous approximations for the inventory and sales functions were not implemented, then a player would not necessarily have a unique best response to the other player's strategy. So this proof of existence would not be valid. Nevertheless, even without the continuous approximation assumption I suspect that a Nash equilibrium would exist for almost all scenarios.□

A Nash equilibrium is a prediction for how the firms will play the game, so it is natural to compare the set of Nash equilibria with the optimal solution. According to the next lemma, the retailer tends to carry less inventory than optimal.

Lemma 6 *If $m>m_r$ and $S_s<S^o$, then $S_r^*(S_s)<S^o-S_s$.*

Proof. If $S_s\geq S^o$, then it is optimal for the retailer to hold zero inventory, and $S_r^*(S_s>$

$S^o) = 0$ is possible, hence the condition $S_s < S^o$. Given $S_s < S^o$, suppose $S_r^*(S_s) \geq S^o - S_s$. Since the retailer's profit function is strictly quasi-concave, that holds only if

$$m_r D'(S^o) \geq h \frac{\partial I_r(S_s, S^o - S_s)}{\partial S_r}. \tag{9}$$

Combine the above with (5),

$$I'(S^o) \geq \left(\frac{m}{m_r}\right) \frac{\partial I_r(S_s^*, S^o - S_s)}{\partial S_r}.$$

But

$$\frac{\partial I_r(S_s, S_r)}{\partial S_r} = I'(S) + \frac{\ln(1/\rho)\rho^{S+1}\left(S_s + 1 - \frac{1-\rho^{S_s+1}}{1-\rho}\right)}{(1-\rho^{S+1})^2} \geq I'(S),$$

which means

$$\frac{\partial I_r(S_s, S^o - S_s)}{\partial S_r} \geq I'(S^o).$$

Since $m > m_r$, it follows that (9) cannot hold. □

The comparable results does not exist for the supplier: the supplier carries too little inventory because it does not receive the full marginal benefit of a sale, $m_s < m$, but the supplier carries too much inventory because it does not incur the holding cost on the additional inventory the retailer must carry when the supplier raises its base stock level, $\partial I_r(S_s, S_r)/\partial S_s > 0$. Either of those effects may dominate.

While the supplier may add too much inventory to the supply chain for a given S_r, according to the next theorem (for all but two special cases) there does not exist a Nash equilibrium in which the supplier's incentive to carry too much inventory compensates for the retailer's incentive to carry too little inventory; the decentralized supply chain performs suboptimally and the decentralized supply chain carries too little inventory relative to the supply chain optimal amount.

Theorem 7 *If $m > m_r$ and $m > m_s$, there does not exist a Nash equilibrium, $\{S_s^*, S_r^*\}$, such that $S_s^* + S_r^* \geq S^o$.*

Proof. For $S_s^* < S^o$, the result follows immediately from Lemma 6. Now consider $S_s^* \geq S^o$. Clearly, in that case $S_r^*(S_s^*) = 0$. Since $m > m_s$, it is not possible that $S_s^* \geq S^o$ because

$$\frac{\partial \pi_s(S^o, 0)}{\partial S_s} = m_s D'(S^o) - hI'(S^o) < mD'(S^o) - hI'(S^o) = \pi'(S^o).$$

□

The decentralized supply chain profit equals the optimal profit only in two uninteresting cases: when one of the player's margin is zero, i.e., when one of the players effectively does not exist.

Lemma 8 *If $m_r = 0$ or $m_s = 0$ then there exists a unique Nash equilibrium, $\{S_s^*, S_r^*\}$, such that $S_s^* + S_r^* = S^o$.*

Proof. If $m_r = 0$, the supplier's first order condition is the same as the supply chain's first order condition, so $S_s^* = S^o$ and $S_r^* = 0$. The analogous argument applies for the retailer when $m_s = 0$.$\square$

These results for the decentralized supply chain are similar to the results found in CZ and CW, but not exactly the same. In both the CZ model and the CW model there is a unique Nash equilibrium, whereas in this model there may be multiple Nash equilibria. But, in all three models, excluding some knife-edge cases (such as when one firms' margin is zero), the optimal solution is not a Nash equilibrium. CZ found that the decentralized supply chain generally carries too little inventory, but sometimes it carries too much inventory. CW also found that the decentralized supply chain sometimes carries too little buffer stock (capacity plus inventory) and sometimes carries too much. In this model the decentralized supply chain always carries too little inventory.

5. Numerical Study

This section provides data from a numerical study that investigates two questions. First, how prevalent are multiple equilibria? Second, what is the magnitude of the competition penalty (the difference in supply chain profit between a Nash equilibrium and the optimal solution as a percentage of the optimal profit)?

All combinations from the following parameters are constructed to form 990 scenarios:

$$\begin{array}{llll} \lambda = \{0.5, 1, 2, 4, ..., 512\} & m_s = \{0.1, 0.2, ..., 0.9\} & b_s = 0 & h = 1 \\ \rho = \{0.05, 0.15, ..., 0.95\} & m_r = 1 - m_s & b_r = 0 & \end{array}$$

The lost sales penalties are set to zero because for fixed margins they do not impact the optimal base stock levels. The remaining parameters are motivated with the optimal base stock condition, (5),

$$\frac{(m_r + m_s)\,\lambda}{h} = \frac{1 - \rho^{S^o+1} - (S^o+1)\rho^{S^o+1}\ln(1/\rho)}{(1-\rho)\rho^{S^o}\ln(1/\rho)}. \tag{10}$$

The left hand side is $m\lambda/h$, so it is sufficient to vary only one of the three parameters; set $m = 1$, set $h = 1$, and let λ vary. However, in the decentralized solution it is important to vary m_r relative to m_s. The supply chain optimal solution should also have some inventory,

$S^o > 0$. Since the right hand side of condition (6) approaches 0.5 as $\rho \to 1$, the left hand side of (6), $m\lambda/h$, should be no less than 0.5, so $\lambda \geq 0.5$. Finally, ρ adjusts the right hand side of the above condition, so the set of ρ values roughly covers the feasible range.

Among the 990 scenarios there are 985 scenarios with a unique Nash equilibrium; scenarios with multiple equilibria are possible, but appear to be rare. Figure 1 displays one of the scenarios with multiple equilibria: $\{m_s = 0.4, \lambda = 128, \rho = 0.95\}$. In the first equilibrium the retailer carries no inventory ($S_s^* = 8.43$, $S_r^* = 0$). The second equilibrium, ($S_s^* = 6.81$, $S_r^* = 1.84$), is not stable; a slight perturbation causes the firms to deviate from the equilibrium, just as a pencil standing on its point is an unstable equilibrium. The third equilibrium, ($S_s^* = 5.14$, $S_r^* = 4.21$), is stable (like the first) and has both firms carrying inventory. It is reasonable to argue that the players will not choose the unstable second equilibrium, but there is no compelling argument to suggest that either of the remaining two equilibria are more likely (or focal) than the other. Thus, in the scenarios with multiple Nash equilibria behavior is less predictable than in the scenarios with only one equilibrium. Fortunately, most scenarios have a unique equilibrium. (If the strategies were restricted to positive integers it is likely that multiple equilibria would be more frequent.)

Table 1 provides statistics on the competition penalty across all scenarios (for the 5 scenarios with multiple equilibria assume the firms coordinate on the equilibrium with the highest total profit, i.e., smallest competition penalty). It is apparent that there are many scenarios in which the competition penalty is relatively small (median = 0.36%), but there are also scenarios with large penalties (average = 10.8%, maximum = 100%). The table clearly indicates that the competition penalty decreases in $m\lambda/h$: as inventory becomes cheaper (h decreases) relative to the maximum profit rate ($m\lambda$) even the decentralized supply chain avoids loss sales and so its profit approaches the supply chain's maximum profit. This is obvious in the extreme case that $h \to 0$: when inventory is free the decentralized supply chain and the centralized supply chain choose sufficiently large base stock levels such that there are essentially no loss sales, and both systems earn $m\lambda$ per unit time.

Table 1: Competition penalty (in percent)

$m\lambda/h$	minimum	median	average	maximum
0.5	2.98	100.00	85.75	100.00
1	0.49	13.02	21.46	97.54
2	0.16	3.81	5.78	23.96
4	0.06	1.54	2.51	10.17
8	0.03	0.70	1.29	5.92
16	0.01	0.34	0.72	5.52
32	0.01	0.19	0.40	3.45
64	0.00	0.11	0.23	2.23
128	0.00	0.05	0.13	0.88
256	0.00	0.03	0.07	0.56
512	0.00	0.02	0.04	0.37
All	0.0004	0.36	10.8	100

Although Table 1 indicates the strong influence of $m\lambda/h$ on the competition penalty, there are other effects at play: when $m\lambda/h = 0.5$ there remains considerable variation in the competition penalty. Figures 2-4 display the competition penalty for different λ, m_r and ρ values. The pattern is apparent: the competition penalty is increasing as the firm's margins become more similar, i.e., as $m_r \to 0.5$; and the competition penalty is increasing as $\rho \to 1$. When $m_r = 0$ or $m_r = m$, one of the firm's margin is the same as the supply chain's margin, so a zero competition penalty results, as documented in Lemma 8. When $m_r = 0.5$, each firm's margin is only 1/2 of the supply chain's margin and so they make poor choices. The pattern for ρ is also explained in terms of the relative impact of holding costs to lost sales: as $\rho \to 0$ lost sales can be avoided with relatively little inventory, whereas a considerable amount of inventory is needed to avoid lost sales as $\rho \to 1$.[1] Overall, an excellent proxy for the competition penalty is the ratio of decentralized supply chain inventory to optimal supply chain inventory, as displayed in Figure 5.

CZ and CW also report that the competition penalty is context specific. They found that the competition penalty is relatively small when the firms incur the same cost of a backorder, but high when the firms have different backorder penalties (i.e., one firm's penalty

[1] The competition penalty is not increasing in ρ when $\rho > 1$. For large values of ρ the supply chain carries very little inventory no matter the base stock levels chosen by the players because the production process is unable to keep up with demand. Hence, the base stock levels chosen by the players do not significantly impact system performance (as long as they are not too low).

is substantially lower than the other). In this model the opportunity cost of a lost sale is analogous to their backorder penalty.

In contrast to both CZ and CW, in this model the competition penalty is low when the firms' lost sales costs are different ($m_r \to 0$ or $m_r \to m$) and high when the firms' lost sales costs are similar ($m_r \to m/2$). There is an explanation for this discrepancy. In the backorder models (CZ and CW) the competition penalty is large not because competition leads to an inappropriate amount of buffer resources (inventory or capacity). For example, in CZ the difference between the decentralized supply chain average inventory and the optimal inventory, as a percentage of the optimal inventory level, is never less than -38%, never more than 4%, and on average only -3.3% in the 2625 scenarios they tested. (In the 990 scenarios tested in this model the average difference is -24%. See Figure 2.) Instead, the competition penalty is large when the allocation of buffer resources is inappropriate: if the retailer carries too little inventory, the supplier is unable to avoid backorders, whereas if the supplier chooses too little of its buffer resource (inventory in CZ, capacity in CW), the retailer must wait longer for its replenishments than it should. Those extreme cases occur when one of the players has no incentive to prevent backorders. In other words, when one player has no incentive to help manage the supply chain (because it does not incur a substantial backorder penalty) the other player is unable to take actions that optimize the supply chain. In this model the allocation of inventory is not important, all that matters is the sum of the firms' base stock levels, $S = S_r + S_s$; even if one firm chooses to hold zero stock the other firm is still able to manage the supply chain optimally by choosing S^o as its base stock level. However, one of the player's incentive to reduce lost sales (its margin, either m_r or m_s) must be similar to the supply chain's incentive to reduce lost sales (its margin, m). When the players have similar margins, neither has the incentive to optimize the supply chain, i.e., double marginalization, Spengler (1950); inventory is too low and the competition penalty is high.

It is likely that there would be a different relationship between the firms' margins and the competition penalty if there were a lead time between the supplier and the retailer in this model. In that case, neither firm could optimize the supply chain if the other firm chose to carry no inventory. That would occur if one of the firms has a small margin; the competition penalty should be significant as $m_r \to 0$ or $m_r \to m$. However, it is still plausible that the competition penalty would be significant as $m_r \to m/2$. Additional research is needed to confirm those conjectures.

There is other supply chain research that finds a context specific competition penalty. Lariviere and Porteus (1999) study a model in which a supplier chooses a wholesale price and a retailer faces a newsvendor problem. They show that the retailer orders less than the supply chain optimal order quantity whenever the supplier earns a positive profit on each unit sold. However, they also show that the competition penalty associated with a simple wholesale price contract decreases as the coefficient of variation of demand decreases. Interestingly, the analogous result holds in this model: the coefficient of variation is decreasing in λ (holding m and h constant), as is the competition penalty.

In Mahajan and van Ryzin (1999) there are competing effects on the retailers' incentive to hold inventory. Since a retailer's margin is less than the supply chain's margin, a retailer tends to carry too little inventory, i.e., double marginalization again. However, when a retailer increases its inventory it lowers the other retailers' demands. Since a retailer ignores that effect on the other retailers, that effect leads the retailers to carry too much inventory. They show that the first effect dominates if there is only one retailer but the two effects become more balanced as the number of retailers increases; the competition penalty decreases as the number of retailers increases.

6. Supply Chain Coordination

The numerical study indicates that the competition penalty is sometime large, sometimes not. When the decentralized supply chain significantly underperforms the centralized supply chain there is an opportunity for the players to divide the spoils from better coordination of their actions. This section studies the kinds of contracts the firms could use to align their incentives to increase the supply chain's profit, where incentive alignment means that the contract changes the firms' profit functions so that the Nash equilibrium in strategies yields the optimal supply chain performance. While the firms could write a contract that directly specifies their actions, i.e., picks base stock levels for each firm, those contracts are not considered for two reasons: first, since each firm will likely prefer to deviate from the contract's specified action, that contract would require a significant amount of monitoring and it would have to impose a substantial penalty for non-compliance; and second, the optimal actions are not independent of the demand rate, so a new contract would be required as the demand rate changes. In short, that contract is too much stick and not enough carrot. Instead, the contracts considered specify a set of payments based on observable and verifiable supply chain metrics: the supplier's average inventory, retailer's average inventory and/or expected lost sales.

The following list provides some of the desirable properties for a coordinating contract: it can coordinate the supply chain in any scenario; both players are better off with the contract than with a wholesale price only contract; the contract allows the players to arbitrarily allocate the gains from incentive alignment; the contract parameters are robust to changes in the demand rate; it is easy to evaluate the optimal contract parameters; and it easy to gather and verify the data required to implement the contract.

6.1 A Retailer Holding Cost Subsidy

From Theorem 7 a coordinating contract must induce the firms to increase supply chain inventory, and, in particular, from Lemma 6 the retailer must carry more inventory whenever $S_s < S^o$. A holding cost subsidy will induce the retailer to carry more inventory and it can temper the supplier's incentive to carry too much inventory: if the supplier subsidizes the retailer's holding cost, the supplier will be less likely to raise its base stock level (which increases the retailer's inventory).

With the holding cost subsidy contract the supplier pays the retailer αh per unit of time per unit of retail inventory, where $0 \leq \alpha \leq 1$: the supplier would never have an incentive to offer $\alpha > 1$, because then the retailer would choose an infinite base stock level; and if $\alpha < 0$, the retailer would certainly choose to carry too little inventory. This contract is analogous to a buy-back contract/return policy (Pasternack, 1985): in effect, the retailer is allowed to instantly return some portion of its unsold inventory to the supplier. Pasternack (1985) demonstrates in a single period (newsvendor) model that a buy-back contract can coordinate the channel (induce the retailer to purchase the optimal order quantity) and it can arbitrarily divide supply chain profits between the two firms; at least in the single period setting, buy-back contracts are quite desirable.

With a holding cost subsidy contract, or α-contract for short, the firms' profits are

$$\pi_r^\alpha(S_s, S_r) = m_r D(S_s + S_r) - (1-\alpha) h I_r(S_s, S_r) - b_r \lambda$$

$$\pi_s^\alpha(S_s, S_r) \;=\; m_s D(S_s+S_r) - hI(S_s+S_r) + (1-\alpha)hI_r(S_s,S_r) - b_s\lambda$$

Let (S_s^α, S_r^α) be a Nash equilibrium given α. It is not difficult to confirm that each player's profit function is unimodal in its strategy and continuous in the other player's strategy, so there does exist a Nash equilibrium. There may even exist an α such that an optimal solution is a Nash equilibrium.

Theorem 9 *If*

$$\frac{1-\rho^{S^o+1}-(S^o+1)\rho^{S^o+1}\ln(1/\rho)}{(1-(S^o+1)\rho^{S^o}+S^o\rho^{S^o+1})\ln(1/\rho)} \le \left(\frac{m_s}{m_r}\right)\left(\frac{\rho}{1-\rho}\right), \tag{11}$$

then there exists a unique α and a unique (S_s^α, S_r^α) pair such that (S_s^α, S_r^α) is a Nash equilibrium and $S_s^\alpha + S_r^\alpha = S^o$.

Proof. As already mentioned, the firms' profit functions are unimodal, so first-order-conditions are sufficient to characterize their optimal responses, which, for the retailer is

$$\begin{aligned} m_r D'(S^o) &= \alpha h \frac{\partial I_r(S_s^\alpha, S_r^\alpha)}{\partial S_r} \\ m_r D'(S^o) &= \alpha h \left(I'(S^o) + D'(S^o)\frac{\rho(1-\rho)}{\lambda}(1-\rho^{S^o+1})I_s(S_s^\alpha,S_r^\alpha)\right) \\ \alpha &= \frac{m_r\lambda}{m\lambda + h\rho(1-\rho)(1-\rho^{S^o+1})I_s(S_s^\alpha,S_r^\alpha)} \end{aligned} \tag{12}$$

For the supplier the condition is

$$\begin{aligned} m_s D'(S^o) &= hI'(S^o) - (1-\alpha)h\frac{\partial I_r(S_s^\alpha,S_r^\alpha)}{\partial S_s} \\ m_s D'(S^o) &= mD'(S^o) - (1-\alpha)h\left(\left(\frac{D'(S^o)}{\lambda}\right)\frac{\rho\left(1-(S_r^\alpha+1)\rho^{S_r^\alpha}+S_r^\alpha\rho^{S_r^\alpha+1}\right)}{\rho^{S_r^\alpha}(1-\rho)^2}\right) \\ m_r\lambda &= (1-\alpha)h\left(\frac{\rho\left(1-(S_r^\alpha+1)\rho^{S_r^\alpha}+S_r^\alpha\rho^{S_r^\alpha+1}\right)}{\rho^{S_r^\alpha}(1-\rho)^2}\right) \end{aligned} \tag{13}$$

The issue is whether there exists an α and a $\{S_s^\alpha, S_r^\alpha\}$ pair that satisfy (12), (13) and $S^o = S_s^\alpha + S_r^\alpha$. Substitute (12) into (13) and rearrange terms,

$$\left(\frac{m_r\lambda}{h}\right)\left(\frac{m\lambda + h\rho(1-\rho)(1-\rho^{S^o+1})I_s(S_s^\alpha,S_r^\alpha)}{m_s\lambda + h\rho(1-\rho)(1-\rho^{S^o+1})I_s(S_s^\alpha,S_r^\alpha)}\right) = \frac{\rho\left(1-(S_r^\alpha+1)\rho^{S_r^\alpha}+S_r^\alpha\rho^{S_r^\alpha+1}\right)}{\rho^{S_r^\alpha}(1-\rho)^2}. \tag{14}$$

The right hand side is strictly convex and increasing in S_r^α. The left hand side is concave and increasing in S_r^α, given that $S_s^\alpha = S^o - S_r^\alpha$. At $S_r^\alpha = 0$, the left hand side is greater than the right hand side (which is zero). Thus, if there exists a solution, it is unique. There exists a solution if the right hand side is greater than the left hand side for $S_r^\alpha = S^o$:

$$\left(\frac{m\lambda}{h}\right)\left(\frac{m_r}{m_s}\right) \le \frac{\rho\left(1-(S^o+1)\rho^{S^o}+S^o\rho^{S^o+1}\right)}{\rho^{S^o}(1-\rho)^2}. \tag{15}$$

Combining (5) with (15) yields (11).□

It is apparent from (11) that an α-contract is more likely to coordinate the supply chain as m_s increases relative to m_r (holding $m_s + m_r$ fixed): if m_s is large relative to m_r then the supplier is likely to be biased to carry too much inventory and so then the α-contract tempers that bias to the point that the supplier chooses $S_s^\alpha = S^o - S_r^\alpha$. The numerical study indicates that (11) is also more likely as ρ increases and as λ increases. Overall, there exists a coordinating α-contract in 282 of the 990 scenarios. Hence, the α-contract is frequently unable to coordinate the supply chain; clearly, a significant limitation.

There are other limitations. Even if a coordinating α-contract exists, it provides only one division of the supply chain profit between the two firms, and there is no guarantee that the α-contract is Pareto improving. Further, the coordinating α-contract must be solved numerically and it almost surely depends on λ.

To summarize, while a holding cost subsidy/buy-back contract is effective in a one period setting, it is significantly limited in a multi-period setting. It is not effective because it is not sufficiently "parameter rich" to coordinate both firms' actions effectively: in the single period setting there is only one action to coordinate (the retailer's order quantity) whereas in the multi-period setting there are two actions to coordinate (the retailer's base stock level and the supplier's base stock level).

6.2 A Lost Sales Transfer Payment

An extra charge for lost sales surely induces a firm to carry more inventory. Consider a lost sales transfer payment contract in which the firms agree to transfer β per expected lost sale per unit time from the supplier to the retailer, where $\beta < 0$ means the retailer pays the supplier for lost sales. With a β-contract the firms' profit functions are

$$\begin{aligned}
\pi_r^\beta(S_s, S_r) &= m_r D(S_s + S_r) - hI_r(S_s, S_r) + \beta(\lambda - D(S_s + S_r)) - b_r\lambda \\
&= (m_r - \beta)D(S_s + S_r) - hI_r(S_s, S_r) + \beta\lambda - b_r\lambda \\
\pi_s^\beta(S_s, S_r) &= m_s D(S_s + S_r) - hI_s(S_s + S_r) - \beta(\lambda - D(S_s + S_r)) - b_s\lambda \\
&= (m_s + \beta)D(S_s + S_r) - hI_s(S_s + S_r) - \beta\lambda - b_s\lambda
\end{aligned}$$

The β-contract adjusts the firms' profit functions in two ways: (1) it adjusts the firms' margins and (2) it adds a fixed component, $\beta\lambda$ or $-\beta\lambda$. Note that the indirect lost sales penalties, b_r and b_s, also provide exactly those two adjustments. The fixed component has no impact on the firms' optimal base stock levels, but clearly the margin adjustments influence the firms' actions. Let $\{S_s^\beta, S_r^\beta\}$ be a Nash equilibrium with a β-contract. There are two β-contracts such that an optimal solution is a Nash equilibrium.

Lemma 10 *If $\beta = m_r$ then the unique Nash Equilibrium is $\{S_s^\beta = S^o, S_r^\beta = 0\}$. If $\beta = -m_s$ then the unique Nash Equilibrium is $\{S_s^\beta = 0, S_r^\beta = S^o\}$. There are no other coordinating β-contracts.*

Proof. Define $\widehat{m}_r = m_r - \beta$ and $\widehat{m}_s = m_s + \beta$. From Lemma 8, the unique Nash Equilibrium has $S_s^\beta + S_r^\beta = S^o$ when $\widehat{m}_r = 0$ or $\widehat{m}_s = 0$, which corresponds to $\beta = m_r$ or $\beta = -m_s$. If $-m_s < \beta < m_r$, then $\widehat{m}_r < m$ and $\widehat{m}_s < m$, so from Theorem 7, there is not an optimal solution that is also a Nash equilibrium. If $\beta > m_r$ ($\beta < -m_s$) then $S_r^\beta = 0$ ($S_s^\beta = 0$)and the supplier's (retailer's) optimal base stock level is greater than S^o.$\square$

If one of those coordinating contracts is implemented, then one of the players earns more than its maximum profit in the supply chain optimal solution without the transfer payment. For example, with $\beta = m_r$ the retailer's profit is $(r - w)\lambda$: the retailer carries no inventory and is fully compensated for every lost sale. Clearly, that is more than the most the retailer could earn in the supply chain optimal solution, $m_r D(S^o)$. If one firm is earning more than its maximum share of the supply chain optimal profit, the other firm must be earning less than its minimum share of the supply chain optimal profit. Given that feature, it will be difficult to get both firms to agree to one of those contracts.

Interestingly, the lost sales transfer contract behaves just like a revenue sharing contract. In a revenue sharing contract the firms agree to divide the supply chain's revenue by some fixed fraction. Adjusting the firms' margins with the β parameter is analogous to adjusting their share of supply chain revenues. Revenue sharing contracts have been successfully implemented in the video rental industry, see Shapiro (1998) and Oestricher (1999). Furthermore, in a single period setting it has many of the same beneficial properties as buy-back contracts (see Pasternack and Drezner, 1999, and Dana and Spier, 1999). However, like the buy-back contract, the revenue sharing/β-contract is not sufficiently parameter rich to be effective in this setting with multiple periods and actions.

6.3 Retailer Holding Cost Subsidy with Lost Sales Transfer

The α-contract's main problem is that it may not be able to coordinate the supply chain. Combining it with a β-contract removes that limitation. To explain, suppose (11) does not hold, so there is no coordinating α-contract. There exists some $\{\overline{m}_r, \underline{m}_s\}$ pair such that $\overline{m}_r + \underline{m}_s = m$ and (11) holds with equality, so it must be that $m_s < \underline{m}_s$. The supplier's adjusted margin with a β-contract is $m_s + \beta$. Hence, there exists a coordinating α/β-contract as long as $\beta \geq (\underline{m}_s - m_s)$.

Adding the β-contract also helps the supply chain with the division of profit, in particular, it shifts profit from the supplier to the retailer. However, it does not always allow the firms to arbitrarily divide the gains from coordination and it may not even be Pareto improving. (Both of those claims are easily verified numerically.) It is true that the firms could use fixed transfer payments to circumvent those limitations: if one firm is earning too little in the contract, a fixed transfer payment clearly alleviates that problem. However, the appropriate fixed transfer payment would depend on λ. Furthermore, there is nothing elegant about resorting to fixed transfer payments. A more refined solution is desirable.

6.4 Inventory and Lost Sales Sharing

With a β-contract the firms adjust their margins, so they are adjusting their share of the lost sales opportunity cost. A similar approach can also be applied to the supply chain's holding cost. Suppose the firms agree to make transfer payments such that the retailer's net holding cost is ϕh per unit of inventory in the supply chain per unit time and the supplier incurs the remaining holding cost. In other words, ϕ is the fraction of the supply chain holding cost assigned to the retailer and $(1 - \phi)$ is the fraction assigned to the supplier. To achieve that outcome the firms transfer from the supplier to the retailer $h\left(I_r(S_s, S_r) - \phi I(S_s, S_r)\right)$ per unit time, where a negative transfer means a payment from the retailer to the supplier. The key feature of this inventory sharing arrangement is that each firm's holding cost per unit time depends only on the sum of the firms' base stock policies (i.e., only on total supply chain inventory) rather than on each firms' own base stock policy (i.e., on how the inventory is allocated between the firms). To provide additional flexibility to this contract, let the firms transfer from the supplier to the retailer β per expected lost sales per unit time, as in the β-contract.

With an inventory and lost sales sharing contract, or ϕ-contract for short, the profit functions are

$$\pi_r^{\phi}(S_s, S_r) = m_r D(S_s + S_r) - \phi h I(S_s + S_r) - b_r \lambda + \beta(\lambda - D(S_s + S_r))$$

$$
\begin{aligned}
&= (m_r - \beta)D(S_s + S_r) - \phi h I(S_s + S_r) - (b_r - \beta)\lambda \\
\pi_s^\phi(S_s, S_r) &= m_s D(S_s + S_r) - (1-\phi)hI(S_s + S_r) - b_s\lambda - \beta(\lambda - D(S_s + S_r)) \\
&= (m_s + \beta)\, D(S_s + S_r) - (1-\phi)hI(S_s + S_r) - (b_s + \beta)\,\lambda
\end{aligned}
$$

Theorem 11 *The ϕ-contract coordinates the supply chain when $\phi = (m_r - \beta)/m$. In that case there is a continuum of Nash equilibria, $S_r^\phi \in [0, S^o]$ and $S_s^\phi = S^o - S_r^\phi$.*

Proof. Begin with the retailer. Express the retailer's profit function as

$$
\pi_r^\phi(S_s, S_r) = \phi\left[\left(\frac{m_r - \beta}{\phi}\right) D(S_s + S_r) - hI(S_s + S_r) - \left(\frac{b_r - \beta}{\phi}\right)\lambda\right].
$$

It is apparent that when $m = (m_r + \beta)/\phi$, which can be written as $\phi = (m_r + \beta)/m$, the retailer's profit function is

$$
\pi_r^\phi(S_s, S_r) = \phi\pi(S_s + S_r) - (b_r + \beta + \phi b)\,\lambda.
$$

Given the above profit function, the retailer clearly will choose $S_r^\phi = \max\{S^o - S_s^\phi, 0\}$, so the retailer's action is optimal. The analogous argument demonstrates that the supplier also chooses the optimal action if $\phi = (m_r - \beta)/m$.$\square$

So why does this contract work? With the lost sales transfer the retailer's share of the supply chain's margin is $(m_r - \beta)/m$. By choosing $\phi = (m_r - \beta)/m$ the retailer incurs the same share of the supply chain's holding cost; i.e., with these adjustments the retailer's and the supplier's profit functions are proportional to the supply chain's profit function. Therefore, both firms want to maximize the supply chain's profit. (Proposition 4 in CW generalizes this intuition.)

The coordinating ϕ-contract is easy to evaluate and independent of the demand rate, λ. However, due to the plethora of Nash equilibria the firms must be sure to coordinate on how they will allocate the inventory in the supply chain, i.e., the firms must ensure that they choose base stock levels that sum to S^o. This coordination should not be difficult, since the firms' profits are independent of the actual division. Nevertheless, the issue cannot be ignored.

The ϕ-contract also provides the firms with the flexibility to allocate the gains from coordination. After some algebra,

$$
\begin{aligned}
\pi_r^\phi(S_s^\phi, S_r^\phi) &= \phi\pi^o + (\phi(r - c) - (r - w))\lambda \\
\pi_s^\phi(S_s^\phi, S_r^\phi) &= (1-\phi)\pi^o - (\phi(r - c) - (r - w))\,\lambda
\end{aligned}
$$

From the above it is apparent that the ϕ parameter allows the firms to arbitrarily divide the supply chain's profit. (If the constraint $\phi \in [0, 1]$ is imposed, then it can be shown that the gains from coordination, π^o minus the wholesale price only Nash equilibrium profit, can be arbitrarily divided, but supply chain profits cannot be in all scenarios.)

While the ϕ-contract is effective and simple, it does require a significant amount of information exchange between the firms: each firm must be able to observe and verify total supply chain inventory and the amount of time the retailer is out of stock. Furthermore, inventory sharing contracts are more difficult to implement when there are multiple retailers: what share of the supplier's inventory is each retailer charged and can each retailer verify that they are charged the correct amount?

7. Discussion

In this supply chain inventory game there exists at least one Nash equilibrium in base stock policies, and, in fact, multiple equilibria may exist. Nevertheless, the numerical study found only a few scenarios with multiple equilibria (5 out of 990 tested scenarios). No matter the number of equilibria, the optimal solution is never a Nash equilibrium; decentralized supply chain inventory management always leads to sub-optimal supply chain performance. In particular, in all Nash equilibria the firms carry less inventory than optimal. But the competition penalty is context specific. (The competition penalty is the difference between the optimal supply chain profit and the Nash equilibrium profit measured as a fraction of the optimal supply chain profit.) The competition penalty is substantial when the holding cost rate is high relative to the potential profit rate (margin per unit times demand rate). The penalty also increases as the firms' margins become more similar: the firms' incentives deviate the most from the supply chain's incentives and so the quality of their decisions deteriorates. Finally, the competition penalty increases in the system's utilization because profit becomes more sensitive to the chosen inventory level as the utilization rate increases: inventory has little impact on total sales when the production rate is high relative to the demand rate, since then lost sales approach zero even if the firms carry little to no inventory.

Several contracts are considered for aligning the firms' incentives. A contract that subsidizes the retailer's holding cost is analogous to a buy-back contract/return policy. While Pasternack (1985) shows that a buy-back contact coordinates the supply chain in a single period setting, in the multi-period setting considered here it is unable to coordinate the supply chain in all scenarios. A lost sales transfer contract, which is analogous to a revenue sharing contract, could coordinate the channel but results in a lopsided allocation of supply chain profit. The combination of the two contracts coordinates the supply chain but restricts the firms' ability to divide the gains from coordination and may not even be Pareto improving.

The most effective contract combines inventory sharing with a lost sale transfer: the retailer incurs a fixed fraction of the supply chain's holding cost, the supplier incurs the remaining fraction and the firms agree to a transfer payment (either from the retailer to the supplier, or vice-a-versa) per expected lost sale. It is shown that there always exists a coordinating contract and the firms are able to arbitrarily divide the gains from coordination.

The primary result from this work is that the decentralized supply chain never performs optimally, but in many scenarios its performance is close to optimal. Since complex contracts are costly to implement (legal fees, monitoring costs, etc.) it is essential that firms confirm that there are significant gains from coordination before embarking on incentive alignment. This may explain why complex contracts are implemented only selectively (e.g., when inventory holding costs are substantial relative to the potential profit rate) and why simple wholesale price only contracts are so prevalent. (Lariviere and Porteus, 1999, make a similar argument based on their single period model.)

This model also identifies why the decentralized supply chain fails to perform optimally; it always stocks less than the optimal amount of inventory. This occurs because inventory in this model is a pure public good: for a given amount of supply chain inventory each firm always prefers to carry less inventory (the sales rate depends only on the total supply chain inventory, so once that is fixed, each firm wants to minimize its own cost). While Cachon and Zipkin (1999) and Caldentey and Wein (1999) study related models (a key distinction is that they allow backorders) in their models buffer resources (inventory in the former, inventory and capacity in the latter) are not purely public goods, i.e., for a fixed amount of buffer resources in the supply chain the firms do not always prefer that the other firm carry more of the resource. (In each model backorders depend on the allocation of buffer resources in the supply chain in addition to the amount of buffer resources.)

Since inventory is a public good in this model, cooperation between the firms *increases* supply chain inventory. That result appears to contradict the substantial amount of anecdotal evidence that cooperation between firms reduces inventory: e.g., Barilla SpA (Harvard case 9-694-046), Procter & Gamble (Harvard case 9-195-126) and H. E. Butt Grocery Company (Harvard case 9-196-061). However, it can be argued that in those cases cooperation expands the set of feasible policies. For example, cooperation may lead to more information sharing which in turn lets the firms implement better policies. That benefit of cooperation is not present in this model because the optimal policy is feasible even if the firms operate independently.

So if firms decide to align their incentives, this research finds that they will probably need a sophisticated contract. It generally is not possible to coordinate the supply chain with a single parameter contract. But even when a single parameter contract can achieve coordination, the contract generally divides supply chain profit in a manner than will not be acceptable to both firms. A parameter rich contract is necessary because there are multiple actions that must be coordinated in the supply chain: the supplier's base stock level and the retailer's base stock level. One suspects that as the complexity of the supply chain model increases, the complexity of the coordinating contract increases as well. It is also worthwhile to note that while a more complex contract may improve the quality of the firms' actions, it also increases the non-trivial costs of information sharing and data verification. Furthermore, while it may be possible to write an intricate contract to align the incentives of two firms, it is significantly harder to write a contract that aligns the incentives of a firm with all of its relationships in a supply chain. Thus, in complex supply chain management there simply does not exist a perfect contract.

Despite that caution on the practical feasibility of sophisticated contracts, there are two reasons why it is nevertheless quite valuable to identify which contracts indeed lead to supply chain coordination. First, identifying the form of a coordinating contact guides firms as to what data they need to collect and verify. In every supply chain there is certainly no lack of data, so developing information systems to collect and track everything is simply silly. Second, the coordination process is facilitated by knowing which contracts lead to overall better performance. This is true because managers are naturally predisposed to assume that the supply chain is a zero sum game: if one firm wins, the other firm must lose. Imagine you are the supplier and the retailer proposes that you partially compensate the retailer for lost sales, i.e., a lost sales transfer payment. You might respond with statements like

“How can I be sure that sales will increase sufficiently to compensate me for this additional cost” or “I can’t give you an incentive to provide poor customer service.” When the firms know that they are not in an exclusively adversarial situation their psychological disposition changes and progress towards supply chain improvement is more likely.

References

Anupindi, R., Y. Bassok and E. Zemel. 1999a. Study of decentralized distribution systems: part 1 - a general framework. Northwestern University working paper.

Anupindi, R., Y. Bassok and E. Zemel. 1999b. Study of decentralized distribution systems: part 2 - applications. Northwestern University working paper.

Cachon, G. 1997. Stock wars: inventory competition in a two echelon supply chain. Fuqua School of Business working paper.

Cachon, G. and P. Zipkin. 1999. Competitive and cooperative inventory policies in a two-stage supply chain. *Management Science.* **45** (7) 936-953.

Caldentey, R. and L. Wein. 1999. Analysis of a production-inventory system. Massachusetts Institute of Technology working paper.

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

Chen, F., A. Federgruen and Y. Zheng. 1997. Coordination mechanisms for decentralized distribution systems with one supplier and multiple retailers. Columbia School of Business working paper.

Dana, J. and K. Spier. 1999. Revenue sharing, demand uncertainty and vertical control of competiting firms. Northwestern University working paper.

Friedman, J. 1986. Game theory with applications to economics. Oxford, Oxford University Press.

Kleinrock, L. 1975. Queuing systems volume I: theory. New York, John Wiley & Sons.

Lariviere, M. and E. Porteus. 1999. Selling to the newsvendor. Fuqua School of Business working paper.

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

Li, L. 1992. The role of inventory in delivery time-competition. *Management Science.* **38** (2) 182-197.

Lippman, S. and K. McCardle. 1995. The competitive newsboy. *Operations Research.* **45** 54-65.

Mahajan, S. and G. van Ryzin. 1999. Inventory competition under dynamic consumer choice.

The Wharton School working paper.

Oestricher, D. 1999. Tales of the tape: tech doesn't scare video rental industry. *Wall Street Journal*: July 30.

Pasternack, B. 1985. Optimal pricing and returns policies for perishable commodities. *Marketing Science*. **4** (2) 166-176.

Pasternack, B. and Z. Drezner. 1999. The videotape rental model. California State University, Fullerton, working paper.

Plambeck, E. and S. Zenios. 1999. Performance-based incentives in a dynamic principal-agent model. Stanford University working paper.

Porteus, E. 1997. Responsibility tokens in supply chain management. Stanford University working paper.

Shapiro, E. 1998. Tired of renting 'shanghai surprise'? Blockbuster, Hollywood hope to help. *Wall Street Journal*, March 25.

Spengler, J. 1950. Vertical integration and antitrust policy. *Journal of Political Economy*. 347-352.

Tayur, S., R. Ganeshan and M. Magazine, Eds. 1998. Quantitative models for supply chain management. Boston, Kluwer.

Tsay, A. 1997. Quantity-flexibility contract and supplier-customer incentives. *Management Science*. **45** (10) 1339-1358.

Wang, Y. and Y. Gerchak. 1999. Supply chain coordination when demand is shelf-space-dependent. University of Waterloo working paper.

Figure 1: Optimal base stock levels when $m_s = 0.4$, $\lambda = 128$, $\rho = 0.95$

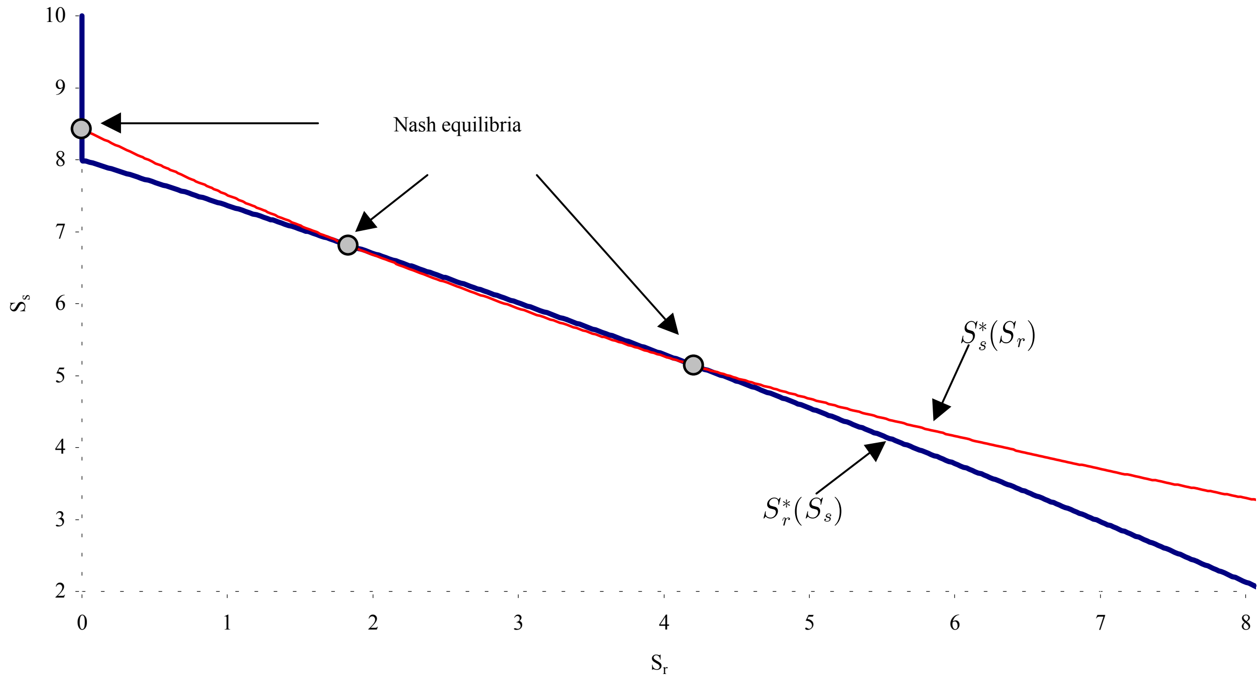

Figure 2: Competition penalty with λ = 0.5

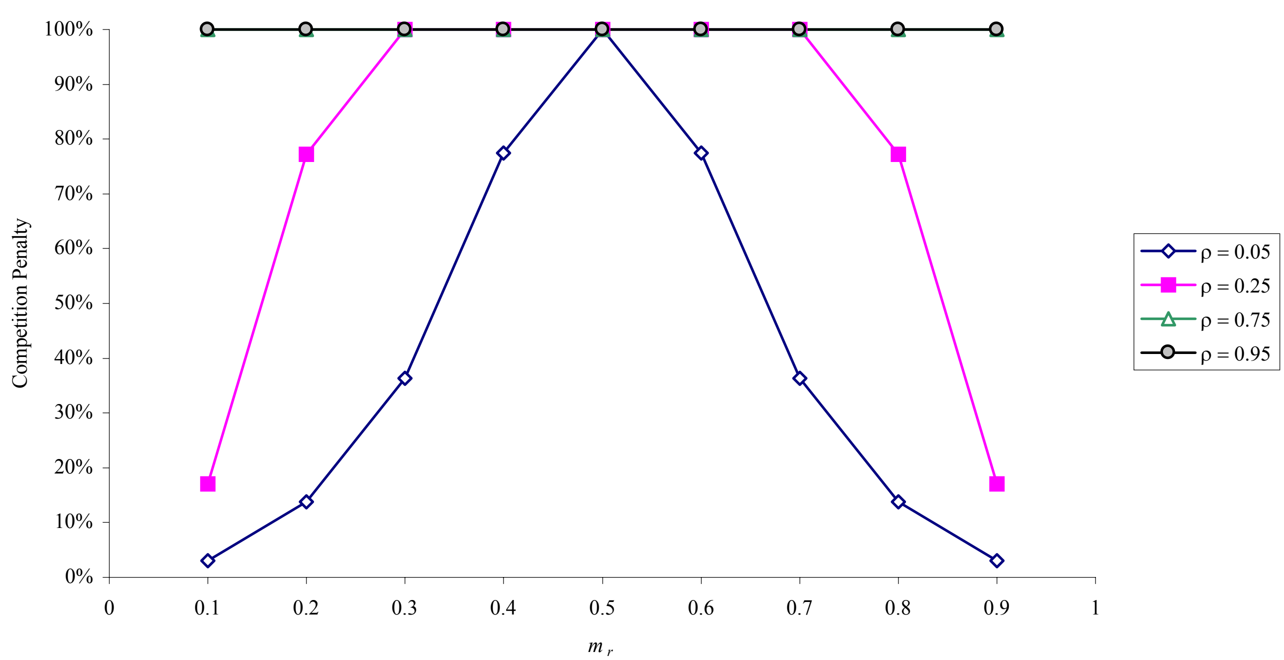

Figure 3: Competition penalty with λ = 1

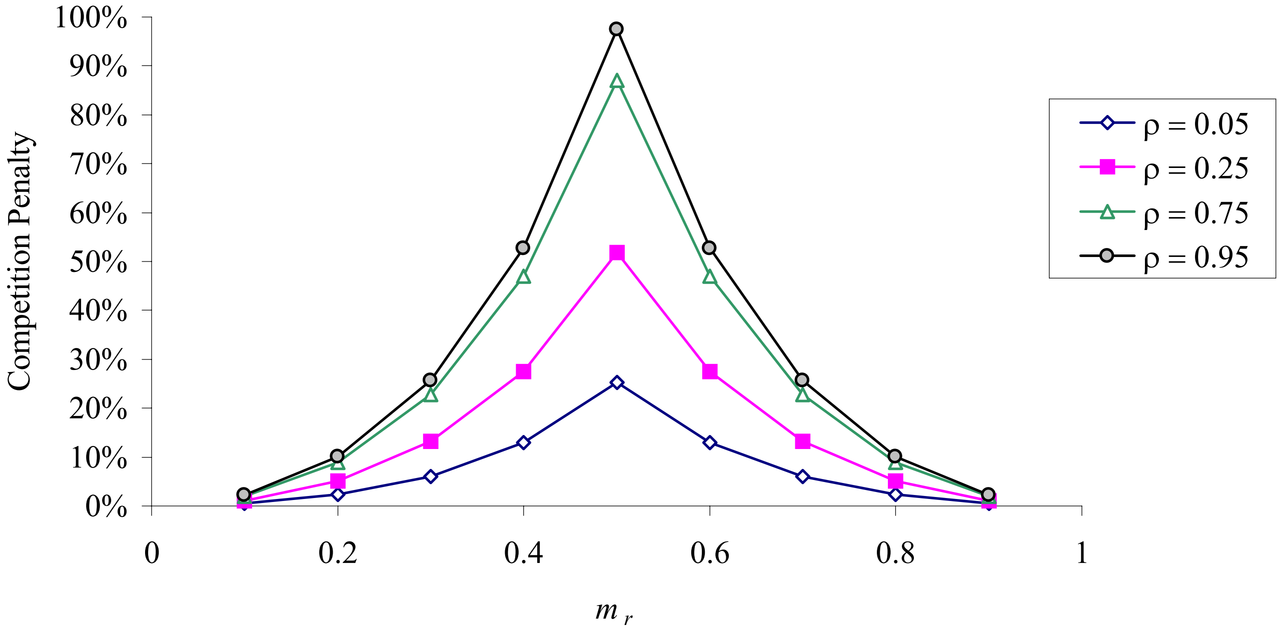

Figure 4: Competition penalty with λ = 2

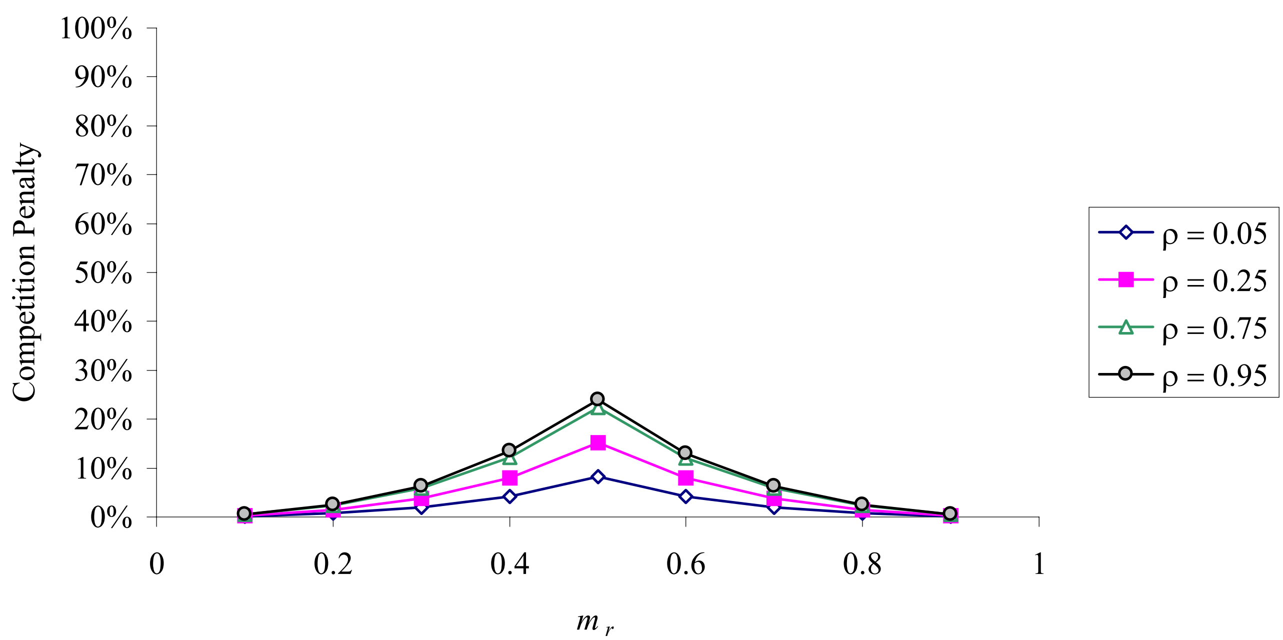

Figure 5: Competition penalty and the ratio of decentralized supply chain inventory to optimal supply chain inventory across 990 scenarios