

1 COMPETITIVE SUPPLY CHAIN INVENTORY MANAGEMENT

Gérard P. Cachon

Fuqua School of Business
Duke University
gpc@mail.duke.edu
www.duke.edu/~gpc

Most supply chains are composed of independent agents with individual preferences. These agents could be distinct firms or they could even be managers within a single firm. In either case, it is expected that no single agent has control over the entire supply chain, and hence no agent has the power to optimize the supply chain. It is also reasonable to assume that each agent will attempt to optimize his own preference, knowing that all of the other agents will do the same. Will this competitive behavior lead the agents to choose policies that optimize overall supply chain performance? The answer is usually “no”, due to supply chain externalities. An externality occurs whenever the action of one agent impacts another agent. For example, suppose agent i 's action benefits agent j . Agent i will tend to do too little of that action because he does not consider the full benefit of the action on the supply chain (assuming increasing that action is costly to agent i). Similarly, suppose agent i 's action confers an additional cost on agent j . In this case agent i will tend to do too much of that action. As will be discussed, many externalities exist in supply chain operations.

When competition degrades supply chain performance the agents can benefit from coordination. But how should they coordinate? They could each agree to work towards the general welfare of the supply chain, but this agreement still leaves each agent with an incentive to serve his own preference. The typical solution is for the agents to agree to a set of transfer payments that modifies their incentives, and hence modifies their behavior. Many types of transfer payments are possible.

This chapter reviews competitive supply chain inventory management. The topic is broadly interpreted, so important research from economics and marketing is included in addition to recent operations management research. §1.1 begins with a model of a supplier that sells a product to a retailer that faces a downward sloping demand curve. In this setting Spengler (1950) obtains the classic double marginalization result: the retailer does not consider the supplier's profit margin when choosing his order quantity, so the retailer orders too little product. Techniques are discussed to encourage the retailer to choose the correct order quantity. §1.2 details a similar model, except the retailer faces stochastic demand at a fixed retail price. Again, the retailer purchases too little inventory because he does not consider the supplier's profit margin. Pasternack (1985) demonstrates that the optimal supply chain profit can be achieved when the supplier offers the retailer a buy-back contract (the supplier purchases unsold goods at a specified buy-back price). §1.3 discusses the use of quantity discounts to raise order quantities. There is a well developed literature on quantity discounts, so this section concentrates on describing the externalities that create the need for quantity discounts.

The primary model in this chapter is discussed in §1.4. This is a two echelon supply chain with a single supplier and a single retailer that faces stochastic demand. The firms incur holding and backorder costs. There are positive lead times between stages and the firms implement bases stock policies. Three externalities are identified in this setting. It is shown that there is a unique Nash equilibrium pair of base stock policies, i.e. there is only one pair of base stock policies such that no firm has a unilateral incentive to deviate. The Nash equilibrium is presumed to be the competitive solution. The optimal policies are never a Nash equilibrium, hence competition always deteriorates supply chain performance. A numerical study indicates that there is substantial variance in the magnitude of the competition penalty: In some cases it is modest (less than 5%), but in other cases it is enormous (over 100%).

Four coordination techniques are applied to this model in §1.5. Cachon and Zipkin (1997) propose linear transfer payments based on actual inventory and backorder levels. Chen (1997) suggests linear transfer payments based on a special accounting of the inventory and backorder levels. Lee and Whang (1996) detail a non-linear transfer payment scheme that is related to Clark and Scarf's (1960) decomposition technique. Porteus (1997) suggests related transfer payments, but implements these payments with responsibility tokens.

Additional operations management research on competitive supply chain inventory management is discussed in §1.6. This section highlights research that identifies additional supply chain inventory management externalities and potential coordination solutions. The final section concludes.

1.1 DOUBLE MARGINALIZATION

Double marginalization is present in most supply chain model because it occurs whenever the supply chain's profits are divided among two or more players and at least one of the players influences demand. To explain in greater detail,

consider a supply chain with a supplier and a retailer that sells a product in one period. The retailer can sell $q \geq 0$ units at price $p(q) \geq 0$. There exists a maximum possible sales quantity \hat{q} , i.e. $p(\hat{q}) = 0$. Over the interval $[0, \hat{q}]$ assume that $p(q)$ is decreasing, continuous, twice differentiable and concave, i.e. $p'(q) < 0$ and $p''(q) \leq 0$. The supplier produces each unit for a cost $c \leq p(0)$, and sells each unit to the retailer for a wholesale price w .

The supplier first announces her wholesale price w and then the retailer chooses an order quantity q . The supplier produces and delivers this order to the retailer. Finally, the retailer sells the q units at price $p(q)$.

To analyze this game, begin with the *centralized solution*, which assumes a single agent controls the entire supply chain to maximize supply chain profits. Next, evaluate the *decentralized solution* which assumes the players make choices with the objective of maximizing their own profits. If the decentralized and centralized solutions differ, investigate how to modify the player's payoffs so that the new decentralized solution corresponds to the centralized solution. This three step analysis pattern is commonplace in the research described in this chapter.

The centrally controlled supply chain's profits are

$$\Pi(q) = q(p(q) - c),$$

which only depends on the retailer's sales quantity. (The supplier's wholesale price decision merely creates a transfer payment between the two firms, so it does not influence supply chain profits.) Since profits are strictly concave in quantity over the interval $[0, \hat{q}]$, the optimal quantity q° satisfies $\Pi'(q^*) = 0$,

$$p(q^\circ) - c + q^\circ p'(q^\circ) = 0. \quad (1.1)$$

(The interval boundaries are never optimal.)

Now assume the retailer must choose a quantity after observing the supplier's wholesale price. The retailer's profits are

$$\pi_r(q) = q(p(q) - w).$$

Since the retailer's profits are strictly concave, his optimal quantity q^* satisfies $\pi_r'(q^*) = 0$,

$$p(q^*) - w + q^* p'(q^*) = 0.$$

The supplier will always choose $w > c$ (otherwise she earns no profit), so it follows from (1.1) that

$$p(q^\circ) - w + q^\circ p'(q^\circ) < 0.$$

Hence, $q^\circ > q^*$. In words, the retailer orders less than the supply chain optimal quantity whenever the supplier earns a positive profits. Spengler (1950) called this problem double marginalization because each firm only considers its own profit margin in making its decision, and does not consider the supply chain's profit margin.

Note that $q^* = q^o$ only when the supplier prices at marginal cost, $w = c$. Hence, *marginal cost pricing* is one solution to double marginalization. Of course, this is not a very good solution for the supplier, since $w = c$ implies she earns a zero profit. In effect, the supplier sells to the retailer her portion of the supply chain for free. A two-part tariff is a better strategy for the supplier. In particular, the supplier could choose to price at marginal cost $w = c$ (part one of the tariff) and also charge a fixed fee $\Pi(q^o)$ (part two of the tariff). The retailer will choose q^o to maximize his gross profits $\Pi(q^o)$, but of course the fixed fee eliminates those profits. Hence, total supply chain profits are optimal and awarded exclusively to the supplier. While marginal cost pricing combined with a fixed fee is a plum strategy for the supplier in this model, Saggi and Vettas (1998) demonstrate that marginal cost pricing is not an effective strategy for the supplier when there are multiple competing retailers.

Without the use of fixed payments, Jeuland and Shugan (1983) suggest the firms coordinate with a *profit sharing contract*. This contract specifies that the supplier earns $f\Pi(q)$, for $0 \leq f \leq 1$, and the retailer earns $(1 - f)\Pi(q)$. The wholesale price is now irrelevant to each firm's profits, so the retailer chooses q^o to maximize his profits, and hence the supply chain earns the optimal profits.

1.2 BUY-BACK CONTRACTS

A buy-back contract specifies a price b at which the supplier will purchase unsold goods from a retailer. These contracts are common in many industries, e.g. publishing and personal computers (see Padmanabhan and Png, 1995).

To understand why a supplier might want to offer a buy-back contract, consider a supply chain with a single supplier and a single retailer. The retailer charges a fixed retail price $p > 0$ and faces stochastic demand. Let $\Phi(x)$ be the cumulative distribution function of demand and let $\phi(x)$ be the density function. Assume $\Phi(x)$ is continuous and differentiable. The sequence of events follows: (1) the supplier announces a wholesale price w and a buy-back price b ; (2) the retailer chooses an order quantity q ; (3) the supplier produces q units at marginal cost c and delivers these units to the retailer ($c < p$); (4) demand is realized and unsold goods are returned to the supplier. Assume the supplier earns nothing from the disposal of the returned units. Pasternack (1985) studies a generalized version of this model.

A central planner chooses only a production quantity to ship to the retailer, since both the wholesale price and the buy-back rate are mere transfer payments. The supply chain's profits are

$$\Pi(q) = -cq + p \left[(1 - \Phi(q))q + \int_0^q x\phi(x)dx \right].$$

The first term is the production cost and the second term is the expected sales revenue. This is a newsvendor problem, and it is well known that the optimal order quantity q^o satisfies

$$\Phi(q^o) = \frac{p - c}{p}. \tag{1.2}$$

The retailer's profits are

$$\pi_r(q) = -wq + p \left[(1 - \Phi(q))q + \int_0^q x\phi(x)dx \right] + b \int_0^q (q - x)\phi(x)dx.$$

The first term is the purchase costs, the second term is expected sales revenue and the third term is expected revenue from returns. Assuming $p > w > b$ (the retailers earns a profit on each unit sold and a loss on each unit returned), the retailer's profits are strictly concave and the optimal order quantity q^* satisfies

$$\Phi(q^*) = \frac{p - w}{p - b}. \quad (1.3)$$

When there is no buy-back, i.e. $w > c$ and $b = 0$, comparison of (1.2) and (1.3) reveals that $q^* < q^o$. In words, if the supplier prices above marginal cost and doesn't offer to purchase unsold goods, double marginalization causes the retailer to order less than the supply chain's optimal quantity. Since supply chain profits depend on q , the sum of the firms' profits will be less than maximum supply chain profits.

The supply chain could achieve its best performance if the supplier were willing to price at marginal cost, but as already mentioned, this is not an attractive solution to the supplier. Instead of lowering her wholesale price to marginal cost, (1.3) indicates that increasing b will raise the retailer's order quantity. In fact, the retailer chooses q^o whenever

$$\frac{p - w}{p - b} = \frac{p - c}{p}. \quad (1.4)$$

Let $\widehat{b}(w)$ be the buy-back that satisfies (1.4),

$$\widehat{b}(w) = p \left(\frac{w - c}{p - c} \right) = \frac{w - c}{\Phi(q^o)}.$$

So supply chain profits are maximized even if $w > c$ as long as $b = \widehat{b}(w)$. Note that $\widehat{b}(w)$ is independent of the demand distribution, so it could apply across multiple retailers facing heterogenous demand distributions.

The supplier's main concern is with her own profits, $\pi_s(w, b, q)$, and not with the supply chain's profits,

$$\pi_s(w, b, q) = q(w - c) - b \int_0^q x\phi(x)dx.$$

The first term is the supplier's sales revenue and the second term is the expected cost of purchasing unsold goods. Assuming the supplier chooses $b = \widehat{b}(w)$, the retailer will choose q^o , so the supplier's profits are

$$\pi_s(w, \widehat{b}(w), q^o) = q^o(w - c) - \widehat{b}(w) \int_0^{q^o} x\phi(x)dx.$$

Differentiate with respect to w

$$\begin{aligned} \frac{\partial \pi_s(w, \widehat{b}(w), q^o)}{\partial w} &= q^o - \frac{p}{p-c} \int_0^{q^o} x \phi(x) dx \\ &= \frac{1}{\Phi(q^o)} \int_0^{q^o} \Phi(x) dx \end{aligned}$$

(The last step is done with integration by parts.) So the supplier's profits are increasing in her wholesale price. In fact, the supplier earns essentially all of the supply chain's profits when $w = p - \varepsilon$, for $\varepsilon \approx 0$. (The retailer's margin approaches zero, but the buy-back rate approaches w , thereby ensuring that the retailer still orders q^o .)

1.2.1 Related models

Several important extensions to Pasternack's model have been considered. Kandel (1996) incorporates asymmetric information, service and quality issues. Emons and Gilbert (1998) relax the assumption of a fixed retail price. Lariviere (1998) provides a more detailed analysis of price-only and buy-back contracts. For instance, he investigates the supplier's optimal wholesale price in a price-only contract.

Padmanabhan and Png (1997) demonstrate that buy-back policies can increase retail competition, thereby benefitting the supplier (even without stochastic demand). Butz (1997), Deneckere, Marvel and Peck (1997), and Deneckere, Marvel and Peck (1996) also study multiple retailer models. They suppose that the retail price depends on the total quantity retailers attempt to sell. While a retailer considers the impact of falling prices on his own inventories, he does not consider how falling prices reduces the value of inventory held by the other retailers. Therefore, after the retailers purchase their inventory they tend to sell this inventory too aggressively, depressing the market price. The retailers can anticipate this behavior when ordering, so they reduce their orders as their expectation for p decreases. In other words, retail competition decreases the incentive to hold inventory because it reduces the retailers' profit margin. The supplier can increase the retailers' orders if she mitigates retail competition. This can be achieved with *resale price maintenance contracts* that set a minimum price that retailers can charge.

Anupindi and Bassok (1998) investigate a model with one supplier and two retailers. Their model's main twist is to incorporate consumer search between the two retailers. They investigate competitive behavior when the retailers make independent decisions as well as when they jointly pool their inventories. In either case, they study how a supplier can benefit from offering a contract with a holding cost subsidy, which is essentially a buy-back contract. (The retailers incur holding costs on end of period inventory, so a holding cost subsidy lets the retailers virtually sell inventory back to the supplier.)

1.3 QUANTITY DISCOUNTS

There is a large literature on quantity discounts so this section concentrates on the reasons why they are used.

Jeuland and Shugan (1983) suggest that quantity discounts can mitigate double marginalization. Suppose the retailer pays $\omega(q)$ for q units, where for quantities less than q^o the marginal price paid is greater than the production cost, i.e. $\omega'(q|q < q^o) > c$, but $\omega'(q^o) = c$. It can be shown that the retailer will choose q^o since his marginal cost at q^o equals the supply chain's marginal cost, c . Further, the supplier earns positive profits since the average wholesale price per unit is greater than c . (See Moorthy, 1987, for additional technical requirements.)

Quantity discounts can also help manage operating costs. Suppose the supplier incurs a fixed order processing cost K_o for each retailer order and let q be a retailer's average order. Thus, the average order processing cost per unit is K_o/q , which is decreasing in q . A retailer does not incur this cost, so a retailer will order a smaller quantity than optimal for the supply chain. Quantity discounts will encourage the retailer to order more.

A retailer's order quantity also can influence the supplier's holding cost. Suppose the supplier incurs a production setup cost, so she will produce in batches. For simplicity, say mq is the supplier's batch size, where m is a positive even integer and q is the retailer's order quantity. Assuming constant retailer demand and a zero lead time between the supplier and the retailer, the supplier's average inventory equals $(m-1)q/2$. Now suppose the retailer doubles his order quantity to $\hat{q} = 2q$. If the supplier doesn't change her batch size, her average inventory is now $(m/2-1)\hat{q}/2$. Since

$$(m/2 - 1)\hat{q}/2 < (m - 1)q/2,$$

the supplier's average inventory has declined. This result holds in more general models. Hence, a retailer tends to order too little because he does not account for the holding cost savings the supplier earns from a larger order quantity. For more extensive treatment of quantity discounts see Lal and Staelin (1984), Lee and Rosenblatt (1986), Weng (1995), and Boyaci and Gallego (1997).

1.4 COMPETITION IN THE SUPPLY CHAIN INVENTORY GAME

This section considers an infinite horizon, stochastic demand inventory game between one supplier and one retailer. The rules of the game are detailed and then the game is compared to other research. The optimal solution is described and the competitive solution is characterized. Finally, several coordination techniques are presented.

1.4.1 Model details

The supplier is stage 2 and the retailer is stage 1. Time is divided into an infinite number of discrete periods. Consumer demand at the retailer is stochastic, independent across periods and stationary. The following is the sequence of

events during a period: (1) shipments arrive at each stage; (2) orders are submitted and shipments are released; (3) consumer demand occurs; (4) holding and backorder penalty costs are charged.

There is a lead time for shipments from the source to the supplier, L_2 , and from the supplier to the retailer, L_1 . Each firm may order any non-negative amount in each period. There is no fixed cost for placing or processing an order. Each firm pays a constant price per unit ordered.

The supplier is charged holding cost h_2 per period for each unit in her stock or on-route to the retailer. The retailer's holding cost is $h_2 + h_1$ per period for each unit in his stock. Assume $h_2 > 0$ and $h_1 \geq 0$.

Unmet demands are backlogged, and all backorders are ultimately filled. Both the retailer and the supplier may incur costs when demand is backordered. The retailer is charged $\alpha_1 p$ for each backorder, and the supplier $\alpha_2 p$, where $\alpha_1 + \alpha_2 = 1$ and $0 \leq \alpha_i \leq 1$. The parameter p is the total system backorder cost, and (α_1, α_2) specifies how this cost is divided between the firms. The parameters (α_1, α_2) are exogenous.

These backorder costs have several interpretations. They may represent the costs of financing receivables, if customers pay only upon the fulfillment of demands. (This requires a discounted-cost model to represent exactly, but the approximation here is standard in the average-cost context, analogous to the treatment of inventory financing costs.) Alternatively, they may be proxies for losses in customer good-will, which in turn lead to long-run declines in demand. Such costs need not affect the firms equally, which is why flexibility is allowed in the choice of $\alpha_i \in [0, 1]$. Finally, they provide an approximation to lost sales.

In period t just before demand define the following for stage i : *in-transit inventory*, IT_{it} , is all inventory in-transit between stages $i + 1$ and stage i ; *inventory level*, IL_{it} , is inventory at stage i minus backorders at stage i (the supplier's backorders are unfilled retailer orders); and *inventory position*, IP_{it} , $IP_{it} = IL_{it} + IT_{it}$. Note that these are local inventory variables and not echelon inventory variables.

Each firm uses a base stock policy: Each period a firm orders a sufficient amount to raise its inventory position plus outstanding orders to that level. Define s_i as stage i 's base stock level. In the inventory game's only move, the players simultaneously choose their strategies, $s_i \in \sigma = [0, S]$, where s_i equals player i 's base stock level, σ is player i 's strategy space and S is a very large constant. (S is sufficiently large that it never constrains the players.) A joint strategy s is a pair (s_1, s_2) . After their choices, the players implement their policies over an infinite horizon. All model parameters are common knowledge (all information is known and verifiable to all players).

Let D^τ denote random total demand over τ periods, and μ^τ denote mean total demand over τ periods. Let ϕ^τ and Φ^τ be the density and distribution functions of demand over τ periods respectively. Assume $\Phi^1(x)$ is continuous, increasing and differentiable for $x \geq 0$, so the same is true of Φ^τ , $\tau > 0$. Furthermore, $\Phi^1(0) = 0$, so positive demand occurs in each period.

This model is identical to the Local Inventory game studied in Cachon and Zipkin (1997). They also consider a model in which firms track echelon inventories and find that the tracking method does influence strategic behavior. The notation in this model is generally consistent with Cachon and Zipkin (1997), but there are some differences. (In Cachon and Zipkin, 1997, an overbar on the local variables distinguishes them from the echelon variables. Echelon variables are not considered, so to avoid notational clutter overbars are not used.)

There are three externalities in this game:

1. The retailer ignores the supplier's backorder costs, so he tends to carry too little inventory;
2. The supplier ignores the retailer's backorder costs, so she tends to carry too little inventory;
3. The supplier ignores the retailer's holding costs, so she tends to carry too much inventory. An increase in the supplier's inventory leads to an increase in the retailer's inventory. (The supplier's average delivery time decreases, thereby raising the retailer's average inventory for a fixed retailer base stock policy.) Higher retail inventory benefits the supplier, through lower backorder costs, but the supplier doesn't pay the retailer's holding costs, so she tends to raise the retailer's inventory more than she should.

Clearly, the second and third externalities conflict so *a priori* it is uncertain whether the supplier will carry too much or too little inventory. While a numerical study confirms that either outcome is possible, it is generally observed that supply chain inventory is too low in the competitive solution.

Chen (1997), Lee and Whang (1996), and Porteus (1997) study similar models. All assume stationary demand, serial supply chains, holding and backorder costs, fixed lead times and common knowledge. Chen (1997) assumes the players attempt to minimize total supply chain costs, so they don't have conflicting incentives. In his model there is a delay between when a stage submits an order and when its upstream supplier receives the order. In this model orders are transmitted instantly. Further, he considers a supply chain with four stages. Both Lee and Whang (1996) and Porteus (1997) study two stage supply chains, but they assume the upstream stage only cares about her local inventory, i.e. $\alpha_2 = 0$. Lee and Whang (1996) assume firms minimize discounted costs instead of average costs.

The supply chain inventory game differs from the models of the previous three sections. Unlike the double marginalization model in §1.1, demand is independent of the player's actions. Unlike the buy-back and quantity discount models, the retailer's orders are not always filled immediately because the supplier may have insufficient inventory.

1.4.2 Optimal Solution

The system optimal solution minimizes the total average cost per period. Clark and Scarf (1960), Federgruen and Zipkin (1984) and Chen and Zheng (1994) demonstrate that a base stock policy is optimal in this setting. The traditional method to find the optimal solution allocates costs to the firms in a particular way. Then, each firm's new cost function is minimized. This section briefly outlines this method.

Let $\widehat{G}_1^o(IL_{1t} - D^1)$ equal the retailer's charge in period t , where

$$\widehat{G}_1^o(x) = h_1[x]^+ + (h_2 + p)[x]^-.$$

(Relative to the retailer's actual costs, the retailer's holding cost is reduced by h_2 per unit, but his backorder penalty cost is increased by h_2 per unit.) Also in period t , define $G_1^o(IP_{1t})$ as the retailer's expected charge in period $t + L_1$, where

$$G_1^o(y) = E \left[\widehat{G}_1^o(y - D^{L_1+1}) \right].$$

Define s_1^o as the value of y that minimizes $G_1^o(y)$:

$$\Phi^{L_1+1}(s_1^o) = \frac{h_2 + p}{h_1 + h_2 + p}. \quad (1.5)$$

This is the retailer's optimal base stock level. Define the induced penalty function,

$$\underline{G}_1^o(y) = G_1^o(\min\{s_1^o, y\}) - G_1^o(s_1^o),$$

and define

$$\widehat{G}_2^o(y) = h_2(y - \mu^1) + \underline{G}_1^o(y).$$

Note that $\underline{G}_1^o(y)$ is non-linear in y .

In period t charge the supplier $G_2^o(IP_{2t})$, where

$$G_2^o(y) = E \left[\widehat{G}_2^o(y + s_1^o - D^{L_2}) \right].$$

The supplier's optimal base stock level, s_2^o , minimizes $G_2^o(\cdot)$.

1.4.3 Game analysis

Define $H_i(s_1, s_2)$ as player i 's expected per period cost when players use base stock levels (s_1, s_2) . The best reply mapping for firm i is a set-valued relationship associating each strategy s_j , $j \neq i$, with a subset of σ according to the following rules:

$$\begin{aligned} r_1(s_2) &= \left\{ s_1 \in \sigma \mid H_1(s_1, s_2) = \min_{x \in \sigma} H_1(x, s_2) \right\} \\ r_2(s_1) &= \left\{ s_2 \in \sigma \mid H_2(s_1, s_2) = \min_{x \in \sigma} H_2(s_1, x) \right\}. \end{aligned}$$

A pure strategy Nash equilibrium is a pair of base stock levels (s_1^*, s_2^*) , such that each player chooses a best reply to the other player's equilibrium base stock level:

$$s_2^* \in r_2(s_1^*) \quad s_1^* \in r_1(s_2^*).$$

Retailer's cost function. In each period, the retailer is charged $h_1 + h_2$ per unit held in inventory and $\alpha_1 p$ per unit backordered. Define $\widehat{G}_1(IL_{1t} - D^1)$ as the sum of these costs in period t ,

$$\widehat{G}_1(y) = (h_1 + h_2)[y]^+ + \alpha_1 p[y]^-.$$

Define $G_1(IP_{1t})$ as the retailer's expected cost in period $t + L_1$,

$$\begin{aligned} G_1(y) &= E \left[\widehat{G}_1(y - D^{L_1+1}) \right] \\ &= (h_1 + h_2) (y - \mu^{L_1+1}) + (h_1 + h_2 + \alpha p) \int_y^\infty (x - y) \phi^{L_1+1}(x) dx. \end{aligned}$$

The retailer's true expected cost depends on both his own base stock as well as the supplier's base stock. After the firms place their orders in period $t - L_2$, the supplier's inventory position equals s_2 . After inventory arrives in period t , the supplier's inventory level equals $s_2 - D^{L_2}$. (The retailer orders D^{L_2} over periods $[t - L_2 + 1, t]$.) When $s_2 - D^{L_2} \geq 0$, the supplier completely fills the retailer's period t order, so $IP_{1t} = s_1$. When $s_2 - D^{L_2} < 0$, the supplier cannot fill all of the retailer's order, and $IP_{1t} = s_1 + s_2 - D^{L_2}$. Hence,

$$\begin{aligned} H_1(s_1, s_2) &= E \left[G_1(\min \{s_1 + s_2 - D^{L_2}, s_1\}) \right] \\ &= \Phi^{L_2}(s_2) G_1(s_1) + \int_{s_2}^\infty \phi^{L_2}(x) G_1(s_1 + s_2 - x) dx. \end{aligned}$$

Supplier's cost function. Define $\widehat{G}_2(IL_{1t} - D^1)$ as the supplier's actual period t backorder cost,

$$\widehat{G}_2(y) = \alpha_2 p [y]^-,$$

and $G_2(IP_{1t})$ as the supplier's expected period $t + L_1$ backorder cost,

$$G_2(y) = E \left[\widehat{G}_2(y - D^{L_1+1}) \right].$$

Define

$$\widehat{H}_2(s_1, x) = h_2 \mu^{L_1} + h_2 [x]^+ + G_2(s_1 + \min \{x, 0\}),$$

so

$$\begin{aligned} H_2(s_1, s_2) &= E \left[\widehat{H}_2(s_1, s_2 - D^{L_2}) \right] \\ &= h_2 \mu^{L_1} + h_2 \int_0^{s_2} (s_2 - x) \phi^{L_2}(x) dx \\ &\quad + \Phi^{L_2}(s_2) G_2(s_1) + \int_{s_2}^\infty \phi^{L_2}(x) G_2(s_1 + s_2 - x) dx. \end{aligned}$$

The first term above is the expected holding cost for the units in-transit to the retailer (from Little's Law), the second term is the expected cost for inventory held at the supplier and the final two terms are the supplier's expected backorder cost.

Equilibrium analysis. The analysis of the game begins by characterizing the cost functions and the best reply mappings.

Theorem 1 $H_2(s_1, s_2)$ is strictly convex in s_2 and $H_1(s_1, s_2)$ is strictly convex in s_1 .

Proof. It is sufficient to demonstrate that the second derivatives are positive. Differentiate $H_2(s_1, s_2)$:

$$\begin{aligned}\frac{\partial H_2(s_1, s_2)}{\partial s_2} &= h_2\Phi^{L_2}(s_2) + \int_{s_2}^{\infty} \phi^{L_2}(x)G_2'(s_1 + s_2 - x)dx \\ \frac{\partial^2 H_2(s_1, s_2)}{\partial s_2^2} &= h_2\phi^{L_2}(s_2) - \phi^{L_2}(s_2)G_2'(s_1) + \int_{s_2}^{\infty} \phi^{L_2}(x)G_2''(s_1 + s_2 - x)dx\end{aligned}$$

The second derivative is positive since $G_2'(y) \leq 0$, $G_2''(y) \geq 0$ and $h_2\phi^{L_2}(s_2) \geq 0$. Differentiate $H_1(s_1, s_2)$,

$$\begin{aligned}\frac{\partial H_1(s_1, s_2)}{\partial s_1} &= \Phi^{L_2}(s_2)G_1'(s_1) + \int_{s_2}^{\infty} \phi^{L_2}(x)G_1'(s_1 + s_2 - x)dx; \quad (1.6) \\ \frac{\partial^2 H_1(s_1, s_2)}{\partial s_1^2} &= \Phi^{L_2}(s_2)G_1''(s_1) + \int_{s_2}^{\infty} \phi^{L_2}(x)G_1''(s_1 + s_2 - x)dx.\end{aligned}$$

Since $G_1(\cdot)$ is strictly convex, $H_1(s_1, s_2)$ is strictly convex in s_1 . \square

Since the cost functions are strictly convex, each player has a unique best reply to the other player's strategy, a useful result to demonstrate existence of an equilibrium. The next two theorems further characterize the best reply mappings. When both players care about backorder costs ($\alpha_1 > 0$, $\alpha_2 > 0$), each player will select a positive base stock. Further, as one player reduces its base stock, the other player will increase its base stock by a lesser amount. This result is used to demonstrate that there exists a unique Nash equilibrium.

Theorem 2 $r_2(s_1)$ is a function; when $\alpha_2 = 0$, $r_2(s_1) = 0$; and when $\alpha_2 > 0$, $r_2(s_1) > 0$ and $-1 < r_2'(s_1) < 0$.

Proof. From Theorem 1 $H_2(s_1, s_2)$ is strictly convex in s_2 , so there a unique base stock that minimizes the supplier's cost. When $\alpha_2 = 0$ the supplier incurs no backorder costs, so she chooses $s_2 = 0$ to incur no holding costs. When $\alpha_2 > 0$, the supplier's first order condition determines her optimal s_2 , but the first order condition is never satisfied at $s_2 = 0$, hence $r_2(s_1) = s_2 > 0$. Assume $s_2 > 0$. From the implicit function theorem,

$$r_2'(s_1) = -\left(\frac{\partial^2 H_2}{\partial s_2 \partial s_1} / \frac{\partial^2 H_2}{\partial s_2^2}\right) \quad (1.7)$$

$$= \frac{-\frac{\partial^2 H_2}{\partial s_2 \partial s_1}}{(h_2 - G_2'(s_1)) \phi^{L_2}(s_2) + \frac{\partial^2 H_2}{\partial s_2 \partial s_1}},$$

where

$$\frac{\partial^2 H_2}{\partial s_2^2} = (h_2 - G_2'(s_1)) \phi^{L_2}(s_2) + \int_{s_2}^{\infty} \phi^{L_2}(x) G_2''(s_1 + s_2 - x) dx,$$

$$\frac{\partial^2 H_2}{\partial s_2 \partial s_1} = \int_{s_2}^{\infty} \phi^{L_2}(x) G_2''(s_1 + s_2 - x) dx.$$

The cross partial of H_2 is positive because $G_2'' > 0$. The denominator of (1.7) is positive because $G_2'(s_1) \leq 0$ and $\phi^{L_2}(s_2) > 0$. So $-1 < r_2'(s_1) < 0$. \square

Theorem 3 $r_1(s_2)$ is a function; when $\alpha_1 = 0$, $r_1(s_2) = 0$; when $1 > \alpha_1 > 0$, $r_1(s_2) > 0$ and $-1 < r_1'(s_2) < 0$.

Proof. $r_1(s_2)$ is a function since H_1 is strictly convex in s_1 . When $\alpha_1 = 0$ the retailer incurs no backorder costs, so he chooses $s_1 = 0$ to incur no holding costs. When $\alpha_1 > 0$, the retailer's first order condition determines his optimal s_1 , but the first order condition is never satisfied at $s_1 = 0$, hence $r_1(s_2) = s_1 > 0$. Assume $1 > \alpha_1 > 0$. From the implicit function theorem

$$\begin{aligned} r_1'(s_2) &= -\frac{\frac{\partial^2 H_1(s_1, s_2)}{\partial s_1 \partial s_2}}{\frac{\partial^2 H_1(s_1, s_2)}{\partial s_1^2}} \\ &= -\frac{\int_{s_2}^{\infty} \phi^{L_2}(x) G_1''(s_1 + s_2 - x) dx}{\Phi^{L_2}(s_2) G_1''(s_1) + \int_{s_2}^{\infty} \phi^{L_2}(x) G_1''(s_1 + s_2 - x) dx}. \end{aligned}$$

From Theorem 2, $s_2 > 0$ (because $\alpha_2 > 0$). Therefore, $\Phi^{L_2}(s_2) G_1''(s_1) > 0$. This implies $-1 < r_1'(s_2) < 0$. \square

Theorem 4 (s_1^*, s_2^*) is the unique Nash equilibrium.

Proof. From Theorem 1.2 in Fudenberg and Tirole (1991), a pure strategy Nash equilibrium exists if (1) each player's strategy space is a nonempty, compact convex subset of a Euclidean space, and (2) player i 's cost function is continuous in s and quasi-convex in s_i . By the assumptions and Theorem 1, these conditions are met, so there is at least one equilibrium. It remains to show that there is a unique equilibrium. Assume $\alpha_1 = 0$; $r_1(s_2) = 0$, and since $r_2(\cdot)$ is a function, $r_2(0)$ is unique. Assume $\alpha_1 = 1$; $r_2(s_1) = 0$, and since $r_1(\cdot)$ is a function, $r_1(0)$ is unique. Now assume $0 < \alpha_1 < 1$. From Theorem 2, $r_2(s_1^*) = s_2^* > 0$. Suppose there are two equilibria, (s_1^*, s_2^*) and (\hat{s}_1, \hat{s}_2) . Assume $s_2^* < \hat{s}_2$. From Theorem 3, this implies that $\hat{s}_1 < s_1^*$. From the same theorem, $r_1'(s_2) > -1$, so $\hat{s}_2 - s_2^* > s_1^* - \hat{s}_1$. But from Theorem 2, $r_2'(s_1) > -1$, which implies $\hat{s}_2 - s_2^* < s_1^* - \hat{s}_1$, a contradiction. The analogous contradiction is obtained if $s_2^* > \hat{s}_2$ is assumed, so the equilibrium is unique. \square

Figures 1 and 2 display the best reply functions in this game as well as the Nash equilibrium and the optimal solution. In neither case does the Nash equilibrium coincide with the optimal solution.

Theorem 5 *Assuming $\alpha_1 < 1$, $s_1^* + s_2^* < s_1^o + s_2^o$.*

Proof. See Cachon and Zipkin (1997), Theorem 15 for proof. \square

It follows immediately from Theorem 5 that the system optimal solution is not a Nash equilibrium whenever $\alpha_1 < 1$. When $\alpha_1 = 1$ the system optimal solution can be a Nash equilibrium under a *very* special condition, see Cachon and Zipkin (1997). Therefore, competitive selection of inventory policies (virtually) always deteriorates supply chain performance (i.e. leads to higher than optimal cost).

A numerical study assesses the magnitude of the competitive penalty (the percentage increase of the Nash equilibrium cost over the optimal cost). The Nash equilibrium and the optimal policies are found for each of the 2625 scenarios constructed from the following parameters:

$$\begin{aligned} \alpha_1 &\in \{0, 0.1, 0.3, 0.5, 0.7, 0.9, 1\} & \alpha_2 &= 1 - \alpha_1 \\ L_1 &\in \{1, 2, 4, 8, 16\} & L_2 &\in \{1, 2, 4, 8, 16\} \\ h_1 &\in \{0.1, 0.3, 0.5, 0.7, 0.9\} & h_2 &= 1 - h_1 \\ p &\in \{1, 5, 25\} \end{aligned}$$

In each scenario demand is normally distributed with mean 1 and standard deviation 1/4. Table 1 summarizes the data.

Table 1: The distribution of competition penalties

α_1	Minimum	5 th percentile	Median	95 th percentile	Maximum
0	107%	117%	804%	5,930%	10,939%
0.1	5%	8%	37%	96%	119%
0.3	2%	3%	9%	19%	26%
0.5	1%	1%	3%	6%	8%
0.7	0%	0%	1%	4%	9%
0.9	0%	0%	1%	17%	45%
1	0%	0%	1%	34%	116%

When each firm cares equally about backorder costs, $\alpha_1 = \alpha_2$, the median competition penalty across scenarios is 3% and the maximum is 8%. When the backorder preferences are asymmetric, $\alpha_1 = 1$ or $\alpha_2 = 1$, then the competition penalty can be quite large. When the retailer cares little about backorders even the median competition penalty is large. This occurs because the retailer chooses to carry little inventory, hence the supplier can do little to prevent consumer backorders. On the other hand, when the supplier cares little about

backorders the median competition penalty is low while the maximum penalty is large. This occurs because the supplier chooses to carry little inventory but in some cases the supplier carries little inventory in the optimal solution. That happens when the supplier's holding cost is large (when $h_1 = 0$ there is no advantage to keeping inventory at the supplier) or when L_2 is small relative to L_1 . Overall, while competition does deteriorate supply chain performance, the magnitude of this problem is clearly context specific.

1.5 COORDINATION IN THE SUPPLY CHAIN INVENTORY GAME

According to Theorem 5, the optimal solution is virtually never a Nash equilibrium. Hence, the firms can lower total costs by cooperatively choosing the optimal base stock levels, (s_1^o, s_2^o) . But any agreement to choose the optimal policy must eliminate each player's incentive to deviate. This is done by appropriately modifying each player's incentives with transfer payments.

1.5.1 Linear contracts

Cachon and Zipkin (1997) propose that the firms adopt a transfer payment contract with constant parameters $(\iota_1, \beta_2, \beta_1)$. This contract specifies that the period t transfer payment from the supplier to the retailer is

$$\iota_1 I_{1t} + \beta_2 B_{2t} + \beta_1 B_{1t},$$

where I_{1t} is the retailer's on-hand inventory, and B_{it} is stage i 's backorders, all measured at the end of the period. There are no *a priori* sign restrictions on these parameters, e.g. $\iota_1 > 0$ represents a holding cost subsidy to the retailer and $\iota_1 < 0$ represents a holding fee. Note that $\iota_1 > 0$ is related to a buy-back contract: having the supplier compensate the retailer for his holding cost is like having the supplier (virtually) purchase some of the inventory back from the retailer.

Define $T_1(IP_{1t})$ as the expected transfer payment in period $t + L_1$ due to retailer inventory and backorders, where

$$\begin{aligned} T_1(y) &= E[\iota_1 [y - D^{L_1+1}]^+ + \beta_1 [y - D^{L_1+1}]^-] \\ &= \iota_1 (y - \mu^{L_1+1}) + (\iota_1 + \beta_1) \int_y^\infty (x - y) \phi^{L_1+1}(x) dx. \end{aligned}$$

Define $T(s_1, s_2)$ as the expected per period transfer payment from the supplier to the retailer,

$$\begin{aligned} T(s_1, s_2) &= E[\beta_2 [s_2 - D^{L_2}]^- + T_1(s_1 + \min\{0, s_2 - D^{L_2}\})] \\ &= \beta_2 \int_{s_2}^\infty \phi^{L_2}(x) (x - s_2) dx + \Phi^{L_2}(s_2) T_1(s_1) \\ &\quad + \int_{s_2}^\infty \phi^{L_2}(x) T_1(s_1 + s_2 - x) dx. \end{aligned}$$

Note that s_1 influences the retailer inventory and backorders, but not the supplier's backorders. Let $H_i^l(s_1, s_2)$ be player i 's costs after accounting for the transfer payment,

$$\begin{aligned} H_1^l(s_1, s_2) &= H_1(s_1, s_2) - T(s_1, s_2), \\ H_2^l(s_1, s_2) &= H_2(s_1, s_2) + T(s_1, s_2). \end{aligned}$$

The objective is to determine the set of contracts, $(\iota_1, \beta_2, \beta_1)$, such that (s_1^o, s_2^o) is a Nash equilibrium for the cost functions $H_i^l(s_1, s_2)$. With these contracts the firms can choose (s_1^o, s_2^o) , thereby minimizing total costs, and also be assured that no player has an incentive to deviate.

To find the desired set of contracts, first assume that H_i^l is strictly convex in s_i , given that player j chooses $s_j^o, j \neq i$. Then determine the contracts in which s_i^o satisfies player i 's first order condition, thereby minimizing player i 's cost. Finally, determine the subset of these contracts that also satisfy the original strict convexity assumption.

The following are the first order conditions:

$$\begin{aligned} \frac{\partial H_1^l}{\partial s_1} &= 0 = \Phi^{L_2}(s_2) (G_1'(s_1) - T_1'(s_1)) \\ &\quad + \int_{s_2}^{\infty} \phi^{L_2}(x) [G_1'(s_1 + s_2 - x) - T_1'(s_1 + s_2 - x)] dx; \end{aligned} \quad (1.8)$$

$$\begin{aligned} \frac{\partial H_2^l}{\partial s_2} &= 0 = -\beta_2 + (h_2 + \beta_2)\Phi^{L_2}(s_2) \\ &\quad + \int_{s_2}^{\infty} \phi^{L_2}(x) [G_2'(s_1 + s_2 - x) + T_1'(s_1 + s_2 - x)] dx. \end{aligned} \quad (1.9)$$

Define $\gamma_2 = \Phi^{L_2}(s_2^o)$. (This is the supplier's in-stock probability, essentially her fill rate.) Furthermore, the supplier's first order condition in the optimal solution is

$$0 = -p + (p + h_2)\Phi^{L_2}(s_2^o) + (h_1 + h_2 + p) \int_{s_2^o}^{\infty} \phi^{L_2}(x)\Phi^{L_1+1}(s_1^o + s_2^o - x)dx,$$

or,

$$\int_{s_2^o}^{\infty} \phi^{L_2}(x)\Phi^{L_1+1}(s_1^o + s_2^o - x)dx = \frac{p - (p + h_2)\gamma_2}{h_1 + h_2 + p}. \quad (1.10)$$

Using (1.8), (1.9) and (1.10) yields the following two equations in three unknowns,

$$\alpha_2 p = \left(\frac{p}{h_1 + h_2} \right) \iota_1 - \beta_1, \quad (1.11)$$

$$h_2 = \left(\frac{h_2}{h_1 + h_2} \right) \iota_1 + \left(\frac{1 - \gamma_2}{\gamma_2} \right) \beta_2. \quad (1.12)$$

It remains to ensure that the costs functions are indeed strictly convex.

Theorem 6 *When the firms choose $(\iota_1, \beta_2, \beta_1)$ to satisfy (1.11) and (1.12), and*

- (i) $h_1 + h_2 > \iota_1 \geq 0$
- (ii) $\beta_2 > 0$
- (iii) $\alpha_1 p > \beta_1 \geq -\alpha_2 p$,

then the optimal policy (s_1^o, s_2^o) is a Nash equilibrium.

Proof. When the following second order conditions are satisfied, H_i^l is strictly convex in s_i , assuming $s_j = s_j^o$, $j \neq i$:

$$\begin{aligned} \frac{\partial^2 H_1^l}{\partial s_1^2} &= \Phi^{L_2}(s_2) (G_1''(s_1) - T_1''(s_1)) \\ &\quad + \int_{s_2}^{\infty} \phi^{L_2}(x) (G_1''(s_1 + s_2 - x) - T_1''(s_1 + s_2 - x)) dx > 0; \\ \frac{\partial^2 H_2^l}{\partial s_2^2} &= (h_2 + \beta_2) \phi^{L_2}(s_2) \\ &\quad + \int_{s_2}^{\infty} \phi^{L_2}(x) (G_2''(s_1 + s_2 - x) + T_1''(s_1 + s_2 - x)) dx > 0. \end{aligned}$$

The first inequality reduces to

$$h_1 + h_2 + \alpha_1 p - \iota_1 - \beta_1 > 0.$$

Substituting (1.11) yields $\iota_1 < h_1 + h_2$ and $\beta_1 < \alpha_1 p$. For the supplier sufficient conditions are

$$\begin{aligned} \alpha_2 p + \iota_1 + \beta_1 &\geq 0; \\ h_2 + \beta_2 + \alpha_2 p + \beta_1 - (\alpha_2 p + \iota_1 + \beta_1) \Phi^{L_1+1}(s_1^o) &> 0. \end{aligned}$$

Combining the first inequality with (1.11) yields $\iota_1 \geq 0$ and $\beta_1 \geq -\alpha_2 p$. The second inequality, along with (1.11) and (1.12), yields $\beta_2 > 0$. \square

These are reasonable conditions: The first requires that the retailer's inventory subsidy not eliminate retailer holding costs; the second stipulates that the supplier be penalized for her backorders; and the third states that the supplier should not fully reimburse the retailer's backorder costs, and the retailer should not overcompensate the supplier's backorder costs.

To help interpret these results, consider the three extreme contracts where one of the parameters is set to zero:

$$\begin{aligned} i) \quad \iota_1 &= 0 & \beta_2 &= \frac{\gamma_2}{1-\gamma_2} h_2 & \beta_1 &= -\alpha_2 p \\ ii) \quad \iota_1 &= h_1 + h_2 & \beta_2 &= 0 & \beta_1 &= \alpha_1 p \\ iii) \quad \iota_1 &= \alpha_2 (h_1 + h_2) & \beta_2 &= \frac{\gamma_2}{1-\gamma_2} \alpha_1 h_2 & \beta_1 &= 0 \end{aligned}$$

(Of these three contracts, the second does not meet the conditions in Theorem 6, because the supplier fully compensates the retailer for all of his costs. The

retailer's incentive to choose the optimal policy is weak: s_1^o is a Nash equilibrium strategy, but any s_1 is too.)

With the first contract the retailer fully reimburses the supplier for the supplier's consumer backorder penalty. However, the supplier still carries inventory because she pays a penalty for her local backorders. With the third contract the supplier subsidizes the retailer's holding costs, but not fully (provided $\alpha_2 < 1$). In addition, the supplier is penalized for her backorders, but less than in the first contract. When the retailer incurs all backorder costs (i.e. $\alpha_1 = 1, \alpha_2 = 0$), only a supplier backorder penalty is required, $\beta_2 = \gamma_2 h_2 / (1 - \gamma_2)$. Hence, the traditional cost allocation scheme used to obtain the optimal solution is needlessly complex. (Recall that scheme requires a retailer holding cost subsidy, a retailer backorder penalty surcharge and a non-linear supplier backorder penalty.)

1.5.2 Accounting inventory

Chen (1997) proposes a linear incentive alignment scheme, but his scheme is not based on actual inventory levels. Instead, his scheme is based on *accounting inventory*, where stage i 's accounting inventory is the actual inventory level it would have if stage $i+1$ always fills its orders immediately. Since the supplier's lead time is assumed to be perfectly reliable, her accounting inventory always equals her actual inventory. On the other hand, the retailer's accounting inventory can be greater than his actual inventory level since the supplier will sometimes only partially fill an order.

To implement accounting inventory to align incentives one of the firms is assigned responsibility for all actual supply chain costs. (In Chen, 1997, this role is played by a single owner who is independent of the managers that actually operate each location.) While either firm can be given this honor, for ease of exposition, assume the supplier bears all costs. This means that each period the supplier pays the retailer's actual holding and backorder costs,

$$(h_1 + h_2) [IL_{1t} - D^1]^+ + \alpha_1 p [IL_{1t} - D^1]^-.$$

This leaves the retailer with zero costs, and hence little incentive to choose the optimal policy. To provide an incentive, each period the retailer pays the supplier h_1^a per accounting inventory, and p^a per accounting backorder. These costs are chosen so that s_1^o minimizes the retailer's payment to the supplier, and hence the retailer will choose s_1^o .

Define

$$\widehat{G}_1^a(x) = h_1^a [x]^+ + p^a [x]^-,$$

which is the retailer's payment to the supplier in a period that ends with accounting inventory x . Define

$$G_1^a(y) = E \left[\widehat{G}_1^a(y - D^{L_1+1}) \right],$$

which is the retailer's expected payment in period $t + L_1$ when he begins period t with accounting inventory y . Since a retailer's accounting inventory always

equals s_1 just before demand (because accounting inventory assumes the supplier always fills the retailer's orders), $G_1^a(s_1)$ is the retailer's expected cost per period. It is easy to confirm that $G_1^a(s_1)$ is strictly convex in s_1 (assuming $h_1^a > 0$ or $p^a > 0$), and the optimal s_1 is

$$\Phi^{L_1+1}(s_1) = \frac{p^a}{h_1^a + p^a}.$$

From (1.5), s_1^o is optimal for the retailer whenever

$$\frac{p^a}{h_1^a + p^a} = \frac{h_2 + p}{h_1 + h_2 + p}. \quad (1.13)$$

The supplier can choose from an infinite number of (\hat{p}^a, \hat{h}_1^a) pairs that satisfy (1.13).

Note the similarity between (1.13) and (1.4). In each case the retailer faces a newsvendor problem so the supplier need only modify the retailer's costs in a manner that sets the retailer's critical fractile equal to the supply chain's optimal critical fractile. In each case the retailer is assured to receive his full order, so the retailer's decision is independent of the supplier's decision, and in turn independent of the demand distribution. In Cachon and Zipkin (1997) the retailer is not assured to receive his full order. Hence, the coordinating parameters depend on the supplier's base stock level and the demand distribution, $\gamma = \Phi^{L_2}(s_2^o)$.

Will the supplier choose s_2^o ? The retailer's payment to the supplier is independent of s_2 , so the supplier does not consider this payment when choosing s_2 . The supplier's actual cost equals the supply chain's cost (because she pays the retailer's actual cost in addition to her own), so (s_1^o, s_2^o) minimizes her actual cost. Hence, the supplier will choose s_2^o and a (\hat{p}^a, \hat{h}_1^a) pair to induce the retailer to choose s_1^o .

While there are many (\hat{p}^a, \hat{h}_1^a) pairs to choose from, the supplier's expected payment from the retailer depends on which pair is chosen. To see this, when the supplier chooses a (\hat{p}^a, \hat{h}_1^a) pair,

$$\begin{aligned} G_1^a(s_1^o) &= \hat{p}^a \left(\frac{1 - \Phi^{L_1+1}(s_1^o)}{\Phi^{L_1+1}(s_1^o)} (s_1^o - \mu^{L_1+1}) \right. \\ &\quad \left. + \frac{1}{\Phi^{L_1+1}(s_1^o)} \int_{s_1^o}^{\infty} \phi^{L_1+1}(x)(x - s_1^o) dx \right). \end{aligned}$$

The term in the above parentheses is independent of (\hat{p}^a, \hat{h}_1^a) , so the supplier can choose to make $G_1^a(s_1^o)$ arbitrarily small or large (by adjusting \hat{p}^a). Hence, the supplier can offer the retailer a contract that leaves the retailer no worse off than in the Nash equilibrium, yet the retailer chooses the optimal policy. This is a great deal for the supplier because the supplier captures all of the benefits of coordination. (Since the retailer's cost is unchanged, the supplier's cost is

reduced by the difference between the Nash equilibrium cost and the optimal cost.)

Notice that the accounting inventory incentive scheme essentially sells the supply chain to the supplier, thereby creating a single agent whose objective is to minimize total supply chain costs. The supplier then modifies the retailer's costs to make the retailer behave as the supplier wishes. Since the supplier buys the supply chain, the supplier can capture all of the benefits of coordination.

1.5.3 Non-linear payments

Lee and Whang (1996) propose an incentive alignment scheme that implements a combination of linear and non-linear transfer payments. Clark and Scarf (1960) also propose these payments to prove that (s_1^o, s_2^o) is optimal, but Lee and Whang demonstrate that with these payments (s_1^o, s_2^o) is a Nash equilibrium.

In the linear portion of this scheme the supplier pays the retailer h_2 for each unit in the retailer's inventory but the retailer pays the supplier $\alpha_2 p + h_2$ per backorder. Let $T_1^n(IL_{1t})$ be the retailer's expected transfer payment to the supplier in period $t + L_1$,

$$T_1^n(y) = -h_2 E[(y - D^{L_1+1})^+] + (\alpha_2 p + h_2) E[(y - D^{L_1+1})^-].$$

Define

$$\begin{aligned} G_1^n(y) &= G_1(y) + T_1^n(y) \\ &= h_1(y - \mu^{L_1+1}) + (h_1 + h_2 + p) \int_y^\infty (x - y) \phi^{L_1+1}(x) dx \end{aligned}$$

so $G_1^n(IL_{1t})$ is the retailer's expected cost in period $t + L_1$.

The non-linear portion of this scheme is a payment from the supplier to the retailer whenever the supplier is unable to fill the retailer's order. Specifically, in period t the supplier pays the retailer $\underline{T}_1^n(s_1, IL_{1t})$,

$$\underline{T}_1^n(s_1, y) = G_1^n(\min\{y, s_1\}) - G_1^n(s_1).$$

When $IL_{1t} = s_1$, the supplier fills the retailer's period t order, so $\underline{T}_1^n(s_1, IL_{1t}) = 0$. When $IL_{1t} < s_1$, the supplier is unable to fill completely the retailer's period t order, so $\underline{T}_1^n(s_1, IL_{1t}) > 0$.

Taking all of the transfer payments into consideration, the retailer's expected per period cost is $H_1^n(s_1, s_2)$,

$$\begin{aligned} H_1^n(s_1, s_2) &= E[G_1^n(\min\{s_1 + s_2 - D^{L_2}, s_1\})] - E[\underline{T}_1^n(s_1, s_1 + s_2 - D^{L_2})] \\ &= G_1^n(s_1) \end{aligned}$$

So the retailer's expected cost is independent of s_2 . Since $G_1^n(y) = G_1^o(y)$, s_1^o minimizes the retailer's cost.

Now consider the supplier's cost, assuming the retailer chooses s_1^o . (This assumption is non-trivial. If $s_1 \neq s_1^o$, the supplier's cost function is not necessarily

convex.) Define

$$\begin{aligned} G_2^n(y) &= E \left[\widehat{G}_2(y - D^{L_1+1}) \right] - T_1^n(y) \\ &= h_2 E[(y - D^{L_1+1})^+] - h_2 E[(y - D^{L_1+1})^-] \\ &= h_2(y - \mu^{L_1+1}) \end{aligned}$$

so $G_2^n(IL_{1t})$ is the supplier's period $t + L_1$ backorder cost minus the retailer's transfer payment. Define

$$\widehat{H}_2^n(s_1, x) = h_2 \mu^{L_1} + h_2[x]^+ + G_2^n(s_1 + \min\{x, 0\}) + \underline{T}_1^n(s_1, s_1 + x)$$

and the supplier's cost including all transfers, $H_2^n(s_1, s_2)$,

$$\begin{aligned} H_2^n(s_1, s_2) &= E \left[\widehat{H}_2^n(s_1, s_2 - D^{L_2}) \right] \\ &= h_2 \mu^{L_1} + h_2 \int_0^{s_2} (s_2 - x) \phi^{L_2}(x) dx + \Phi^{L_2}(s_2) G_2^n(s_1) \\ &\quad + \int_{s_2}^{\infty} \phi^{L_2}(x) G_2^n(s_1 + s_2 - x) dx \\ &\quad + \int_{s_2}^{\infty} \phi^{L_2}(x) G_1^n(s_1 + s_2 - x) dx - (1 - \Phi^{L_2}(s_2)) G_1^n(s_1) \end{aligned}$$

Simplify the above,

$$\begin{aligned} H_2^n(s_1, s_2) &= h_2(s_2 + s_1 - \mu^{L_2+1}) + \int_{s_2}^{\infty} \phi^{L_2}(x) G_1^n(s_1 + s_2 - x) dx \\ &\quad - (1 - \Phi^{L_2}(s_2)) G_1^n(s_1). \end{aligned}$$

It is easy to show that $H_2^n(s_1^o, s_2) = G_2^o(s_2)$, so $H_2^n(s_1^o, s_2)$ is minimized with $s_2 = s_2^o$. Therefore, (s_1^o, s_2^o) is a Nash equilibrium since each player minimizes its cost given the strategy choice of the other player.

1.5.4 Responsibility tokens

Porteus (1997) proposes an incentive scheme, called responsibility tokens, that blends some of the features of the non-linear scheme in Lee and Whang (1996) with some of the features of accounting inventory. As with accounting inventory, the retailer receives perfectly reliable deliveries from the supplier, but the implementation of this assumption is different. Whenever the supplier is unable to fill a retailer order she issues responsibility tokens to cover the portion of the order she cannot fill. From the retailer's perspective these responsibility tokens are equivalent to real inventory. Tokens issued in period t are received in period $t + L_1$. The retailer incurs holding costs on these tokens and can use them to prevent backorder costs. To explain, suppose τ tokens are issued in period t and y is the retailer's actual inventory level in period $t + L_1$ when costs are measured. The retailer's actual holding and backorder costs are

$$(h_1 + h_2)[y]^+ + \alpha_1 p[y]^-,$$

but the supplier pays the retailer $\alpha_1 p \min\{[y]^-, \tau\}$ to compensate the retailer for those backorder charges that the retailer could have avoided if his tokens were real inventory. However, the retailer pays the supplier $(h_1 + h_2)[\tau - [y]^-]^+$ to reward the supplier for saving the retailer actual holding costs. Once the retailer's actual costs are combined with these transfer payments, the retailer's costs are

$$(h_1 + h_2)[y + \tau]^+ + \alpha_1 p [y + \tau]^-. \quad (1.14)$$

There are some additional transfer payments. As in Lee and Whang (1996), the supplier pays the retailer h_2 per unit of inventory and the retailer pays the supplier $\alpha_2 p + h_2$ per backorder, where inventory and backorders include tokens. Following the above example, the retailer receives $h_2[y + \tau]^+$ from the supplier but pays $(\alpha_2 p + h_2)[y + \tau]^-$ to the supplier. Combining these transfers with (1.14), the retailer's costs in period $t + L_1$ are

$$h_1[y + \tau]^+ + (p + h_2)[y + \tau]^-$$

Since $y + \tau = s_1 - D^{L_1+1}$, $G_1^n(s_1)$ is the retailer's expected cost per period, which is minimized by s_1^o .

Since s_1^o minimizes the retailer's actual costs plus his transfer payments for any s_2 , the retailer will choose s_1^o . Given that $G_1^n(s_1^o)$ will be the retailer's cost, the supplier's costs must equal all remaining costs in the system, which equals $G_2^o(s_2)$. Hence, the supplier also will choose the optimal policy, s_2^o , and (s_1^o, s_2^o) is a Nash equilibrium.

Although there is a strong resemblance between responsibility tokens and the non-linear scheme in Lee and Whang (1996), there is a subtle difference. With responsibility tokens the retailer receives perfectly reliable supply, so the retailer's costs are independent of the supplier's base stock. (This is also the case with accounting inventory.) Hence, it is the supplier that bears the actual cost consequence of her late deliveries. With the non-linear scheme the retailer doesn't receive perfectly reliable supply, but the supplier does pay the retailer an amount that exactly compensated the retailer for the expected cost consequence of any late delivery. Hence, in this case the retailer pays the actual cost consequence of late deliveries. This distinction is immaterial when all of the players are risk neutral (as is assumed in each case). However, it becomes relevant if one of the players were risk averse. For example, if the retailer were risk averse then the supply chain may be better off using responsibility tokens to let the risk neutral supplier bear the consequence of late deliveries. On the other hand, if the supplier were risk averse, the supply chain may be better off using the non-linear scheme since then the supplier only pays the expected consequence of her action. (In general, the most risk neutral player should bear the supply chain's risk. See Tirole, 1990, for a discussion of risk sharing among players with heterogenous tastes for risk.)

1.6 OTHER RESEARCH

There are many other supply chain inventory management models that differ significantly from the ones already described. The section reviews these models and highlights additional supply chain externalities and coordination techniques.

1.6.1 Multiple Retailers with Stochastic Demand

Multiple retailers substantially complicates supply chain inventory analysis. Optimal policies are not known and it is difficult even to find the best policy within a reasonable class of policies, e.g. reorder point policies. Nevertheless, some results have been obtained.

Cachon (1997) considers a two echelon supply chain with a single supplier and N retailers. Each retailer faces identically distributed Poisson demand. Inventory is review continuously and each location implements an (R, Q) policy: when a location's inventory position equals R it orders Q units from its supplier. There is a fixed lead time to replenish the supplier and a fixed lead time between the supplier and each retailer. There are holding costs and both the supplier and the retailers care about consumer backorder costs, as in §1.4.

In this model Axsater (1993) demonstrates how to find optimal (R, Q) policies assuming centralized control. Cachon (1997) assumes each location selects its reorder point to minimize its own costs given that each other player will do the same. The search for Nash equilibria in reorder point policies is complex because discrete demand implies that each player's strategy space is discrete (i.e. non-convex), which precludes the implementation of standard existence theorems. (It is common that there does not exist an equilibrium in games with non-convex strategy spaces.) However, it is shown that this is a supermodular game. Roughly speaking, in a supermodular game each player's strategy space is ordered and as player j chooses a higher strategy, player i will also wish to choose a higher strategy. In the inventory game the supplier's strategy space (reorder points) are given the natural ordering, i.e. a higher reorder point is a higher strategy, but the retailers are given the opposite ordering, i.e. a higher reorder point is a lower strategy. Hence, as the supplier increases her reorder point, a retailer will decrease his reorder point. See Milgrom and Roberts (1990) for a review of supermodular games. The supermodular property implies that there exists a pure strategy Nash equilibrium, and further, it provides an algorithm to find all Nash equilibria.

In a numerical study it is found that there can be multiple equilibria and the optimal reorder point policies can be a Nash equilibrium. These findings contrast with the results from the single retailer model. However, as in the single retailer model, the competition penalty is relatively moderate when the firms incur equal backorder costs but often quite large when the firms incur asymmetric backorder costs. Additional research is needed to better understand why these results occur. There are several candidate explanations: the multi-retailer model has a discrete strategy space, whereas the single retailer model

has a continuous strategy space; and in the multi-retailer model the allocation of inventory among the retailers is an issue, whereas it is not in the single retailer model.

Anderson, Axsäter and Marklund (1996) study almost the same model, but they do not investigate competitive decision making. Instead, they provide a mechanism to decouple the supply chain: each location solves a problem that is independent of the reorder points chosen by the other locations. Specifically, the mechanism imposes on the supplier a linear backorder penalty cost. The supplier's problem is to choose a reorder point that minimizes her actual holding costs plus this backorder penalty cost. This problem depends on the retailers' ordering process, but the retailers' ordering process is independent of their reorder points. So the supplier can solve her problem independent of the retailers. The retailers minimize their actual holding and backorder costs assuming a deterministic lead time which equals the retailers' expected actual lead time. Anderson, et al. (1996) are unable to explicitly determine the correct backorder penalty, so they search for a good penalty via an iterative procedure. The procedure is not guaranteed to converge to the optimal reorder points, but they found that the procedure worked quite well in a numerical study.

In a significantly different model, Hausman and Erkip (1994) also study a decoupling technique to coordinate a supply chain model originally studied by Muckstadt and Thomas (1980). This model has multiple retail locations, multiple products and emergency shipments. Instead of transfer payments, they impose fill rate constraints on each location and show how to choose these fill rates so that each location selects near optimal policies.

1.6.2 Multiple Retailers with Deterministic Demand

Chen, Federgruen and Zheng (1997) consider a multi-echelon supply chain with multiple retailers. Each retailer faces deterministic demand that is decreasing in the retailer's price. Demand curves can differ across retailers. Firms incur holding and ordering costs. To manage each retailer's account, the supplier incurs a fee that is increasing and concave in the retailer's purchase quantity. Hence, the account management fee per unit sold is decreasing in sales. Retailers choose their retail prices and their order intervals (time between orders, which equals the order quantity divided by the demand rate). The supplier chooses her wholesale price and her order interval. They assume all of the firms choose power-of-two policies: a base order interval is chosen and then all firms choose their order interval to be a power of two multiple of this base interval. For fixed prices (and hence demand rates), Roundy (1985) demonstrates that power-of-two policies are guaranteed to be within 2% of the optimal cost.

Several incentive conflicts occur in this model. The retailers tend to price too high because they don't consider the supplier's profit when choosing their sales rate, i.e. double marginalization. They are also biased towards lower than optimal sales because they don't incur the account management fee. This externality is somewhat different than the order processing externality described in §1.3 because it is based on average sales and not the order frequency. Finally,

the supplier's holding cost is non-increasing in the retailer's order interval, so the retailers tend to choose order intervals that are shorter than optimal.

Unlike in the stochastic demand models, in the deterministic model it is implicitly assumed that the supplier never makes a late delivery. Hence, the supplier's order interval has no impact on the retailers' costs. Only the supplier's wholesale price affects the retailers' costs.

Chen, et al. (1997) demonstrate that when the supplier chooses the optimal order interval, the supplier can announce a wholesale price schedule that makes the retailers choose the optimal retail prices and order intervals. This wholesale price per unit equals the supplier's marginal cost plus a sales volume fee plus an order interval fee. The sales volume fee equals the supplier's account management cost per unit given a retailer's sale volume, so the retailers exactly compensate the supplier for her account management costs. The order interval fee is non-increasing in a retailer's order interval length, and exactly compensates the supplier for her holding costs. With this scheme the supplier incurs her ordering costs, but the retailers incur all other costs, either directly or indirectly through the wholesale price. Further, the retailers receive all sales revenues, so they earn more than the supply chain's optimal profits. To gain the supplier's participation, the retailers clearly must share some of the profits via lump sum payments.

1.6.3 Capacity Allocation

Consider a supply chain with a single supplier and multiple retailers. Demand occurs in a single period. The supplier has limited capacity, so whenever the retailers order more than capacity, the capacity must be allocated among the retailers. How much will each retailer order and how much capacity will the supplier build?

Lee, Padmanabhan and Whang (1997) call this the *shortage game*. In their version of the game the supplier has uncertain yield, so even though the retailers know how much they all want, they don't know how much will be available. If they order more than the supplier produces, the supplier allocates the production proportional to their order quantities. Lee, et al. (1997) demonstrate that each retailer will order more than their desired quantity.

Cachon and Lariviere (1996) study a different shortage game. The supplier's capacity is known but each retailer has private information about his own demand: while a retailer knows his own expectation of demand, he does not know the demand expectations of the other retailers. Hence, each retailer is uncertain about how much the other retailers will order. There are several incentive problems in this setting. The retailers tend to purchase too little relative to the supply chain optimal, because the supplier charges above marginal cost, i.e. double marginalization. But when capacity is scarce, a retailer tends to order too much because he does not consider the value of scarce inventory to other retailers: the supply chain should allocate inventory to the retailer with the highest marginal value for inventory, but a retailer with a low marginal value may nevertheless receive a larger allocation by increasing his order. There are

also incentive problems with the supplier. Since the supplier faces random demand (she doesn't know the retailers' expectations), and since she does not consider the retailers' profit margin on each unit, the supplier tends to build less than the supply chain optimal capacity.

Cachon and Lariviere (1996) identify a broad class of allocation schemes that induces truth telling behavior among the retailers, i.e. the retailers order their desired quantities. Another class of allocation schemes induces the retailers to inflate their orders. In a numerical study it is found that the order inflation schemes generally outperform the truth-inducing schemes. There are several reasons for this result. Double marginalization makes the retailers order too little, but an order inflation allocation scheme mitigates this problem. The problem with order inflation is that it might lead to a poor allocation of inventory among the retailers. The worst case scenario is that a retailer with low demand ends the period with excess inventory while a retailer with high demand ends the period with substantial shortages. However, this generally doesn't happen in equilibrium, because equilibrium orders correspond with needs, i.e. the retailer with the highest need inflates his order the most, thereby receiving the greatest allocation. So even though retailers inflate their orders, inventory allocations tend to be reasonable. Finally, retailer order inflation also induces the supplier to build more capacity, which can benefit everyone in the supply chain.

Several other shortage gaming models have been studied. In a multi-period model, Cachon and Lariviere (1997) study turn-and-earn allocation of scarce capacity, a scheme commonly used in the automobile industry. Mallik and Harker (1997) study capacity allocation in an internal supply chain with multiple production managers. See Ha (1996) and Corbett (1998) for other supply chain inventory models with asymmetric information.

1.6.4 Information and Production Timing

Predicting demand for seasonal products is notoriously difficult. However, it has been frequently observed that early season sales provide excellent indications of total season sales (see Fisher and Raman, 1996). So retailers would like to be able to replenish inventory after observing early season sales. But delayed production is usually more expensive than early production, so suppliers prefer to have full production commitments well before the season begins. Hence, firms face a conflict between early cheap production with unpredictable demand and late expensive production with reliable demand.

Donohue (1996) studies this problem in a supply chain with one manufacturer and one buyer. The buyer places an initial order and then the manufacturer decides her initial production quantity. The buyer then observes a signal that improves the buyer's demand forecast. The buyer places a second order and the manufacturer decides her second production quantity. This second production is more expensive than the first. Finally, actual demand is realized. She demonstrates that a buy-back contract makes the firms choose the centrally

optimal decisions. Further, the buy-back price is independent of the demand distributions, as in §1.2.

Barnes-Shuster, Bassok and Anupindi (1998) expand upon Donohue's model. In their model there are two demand periods. The sequence of events follows: the buyer makes firm orders for delivery in each period and also purchases options for additional units to be delivered after the first demand period; the supplier purchases raw material and makes an initial production; the buyer receives his first period order and stochastic demand occurs; the buyer receives his period two firm order plus he can exercise his options for a per unit exercise fee; the supplier delivers the buyer's period two firm order plus exercised options; and stochastic demand occurs. The buyer incurs a per unit holding cost for inventory carried from period one to two. The supplier's total production over the two periods is limited by the quantity of raw material purchased at the start of the game. Hence, the supplier is required to purchase sufficient raw material to cover the buyer's firm orders and potentially all of his options. In addition to the per unit cost of raw material, there is a per unit production cost which is higher in period two than in period one.

They find that marginal cost pricing of these options contracts doesn't coordinate the channel. The issue is the period two option exercise price. Since period one production is cheaper than period two production, and since there is an inventory holding cost at the retailer, the supplier may have more inventory at the start of period two than is needed to fill the buyer's period two firm order. From the supply chain's perspective, these excess units should certainly be moved to the retailer because they cannot fill period two demand while at the supplier. Only with a zero exercise price will the buyer necessarily order these additional units. However, a zero exercise price also means that the buyer will necessarily exercise all of his options. This may require the supplier to conduct some period two production, which has a greater than zero marginal cost. Hence, a zero exercise price may cause the buyer to exercise too many options from the supply chain's perspective, but a non-zero option may cause the buyer to exercise too few options. To summarize, marginal cost pricing doesn't work in this setting because the supply chain's marginal cost of moving additional units to the buyer in period two is not independent of the quantity moved.

Eppen and Iyer (1997) also study option like contracts, which they call backup agreements. However, they do not consider the issue of channel coordination.

Tsay (1997) studies a setting in which the supplier makes one production, which occurs after the retailer's initial order, but before the retailer observes a signal of demand. The retailer places his firm order after observing this signal. The firm order may or may not be constrained by the initial order. If there is no link between them, the retailer is likely to provide a meaninglessly high initial order. Tsay (1997) proposes to link them via a *quantity flexibility contract*. If q is the initial order, these contracts specify that the firm order can be no greater than $(1 + \alpha)q$ and no less than $(1 - \omega)q$, for $\omega \in [0, 1]$ and $\alpha \geq -\omega$. Since the supplier allows the retailer to request up to $(1 + \alpha)q$, Tsay assumes the

supplier produces $(1 + \alpha)q$, thereby guaranteeing that the retailer could indeed receive his maximum order. (See Cachon and Lariviere, 1997, for a discussion of a model in which it is not assumed that the supplier will necessarily live up to her commitments.) It is shown that these contracts can guarantee that each firm will choose supply chain optimal actions. Hence, $(1 + \alpha)q$ must equal the supply chain optimal production quantity.

Quantity flexibility contracts are more complex than buy-back contracts. A buy-back contract allows a retailer to return any portion of his initial order at a buy-back price which is less than the wholesale price. A quantity flexibility contract allows the firm to return a limited portion of his initial order for a full refund of the wholesale price. Further, a buy-back contract delivers to the retailer his initial order without an opportunity to increase the order. This makes sense when the retailer doesn't receive additional information between the initial order and the delivery (as is assumed in §1.2). However, because in Tsay's model the retailer does receive a demand signal, the quantity flexibility contract allows the retailer to increase his order somewhat. See Tsay and Lovejoy (1995) for additional discussion on the implementation of quantity flexibility contracts.

1.6.5 Internal Markets

Porteus and Whang (1991) consider a "supply chain" model with one owner, one manufacturing manager and several product managers. The manufacturing manager's costly effort affects the realization of capacity, which is subject to a random shock. The product managers' costly effort affects the realization of sales, which is also random. The owner wishes to maximize her profits, subject to the conditions that she cannot choose the effort levels of her managers. Further, she must offer them a sufficiently attractive contract to prevent them from seeking employment at another firm. Porteus and Whang (1991) suggest that the owner can achieve her objective through an internal market. The product managers receive all revenues from the sale of their product and pay the manufacturing manager the realized marginal value of capacity. (The marginal value of capacity is decreasing in the realization of capacity.) In addition, the product managers pay a fixed fee that equals their expected profits, so they are indifferent between working for the owner and their best outside opportunity. The manufacturing manager receives the expected marginal value of each unit of capacity she delivers, but also pays a fixed franchise fee to the owner that equals the manufacturing manager's expected profits. Interestingly, the owner loses money on average when operating this market: The average price the marketing managers pay for capacity is less than the price the manufacturing manager receives. Hence, the owner earns her profits through the franchise fees, and operates the internal market merely to induce optimal behavior. Kouvelis and Lariviere (1997) generalize the concept of an internal market, and apply it to several other settings.

1.6.6 *The Beer Game*

The Beer Game, Sterman (1989), is probably the most famous demonstration that decentralized decision making can lead to poor supply chain performance. There are several explanations for this result. Of the four players only the retailer observes demand, so it is difficult for upstream locations to know how much they should order. Indeed, optimal policies for this setting are not known, even if all of the players were to observe demand. Furthermore, poor judgment may also be a culprit: there is evidence that players forget to consider previously ordered inventory when choosing order quantities, leading them to order too much. Incentive conflicts are notably absent from this list of explanations: the players are told to minimize total supply chain costs.

1.6.7 *Vendor Managed Inventory*

With Vendor Managed Inventory (VMI) a supplier assumes responsibility for choosing a retailer's stocking level. In exchange for control of supply chain inventory, the supplier agrees to charge the retailer a constant wholesale price. Since the supplier determines shipments to the retailers, VMI generally requires electronic transmission of inventory and demand data from the retailer to the supplier, e.g. Electronic Data Interchange (EDI). The grocery industry provides several examples of successful VMI implementation. (These systems are also called Continuous Replenishment Programs, CRP, or Continuous Product Replenishment programs, CPR.)

Clark and Hammond (1997) studied retailers that adopted VMI with the Campbell Soup Company. The performance of these retailers was compared with retailers that implemented EDI but did not transfer control of inventory to Campbell Soup, i.e. they merely used EDI to submit orders electronically. They found that the VMI retailers experienced substantially better performance gains over the latter group, suggesting that transferring control provided significant benefits. Cachon and Fisher (1997) also studied Campbell Soup's implementation of VMI. However, they found that operating benefits could have been achieved even if the retailers had maintained control of their own inventories.

Several models have been developed to study VMI. Cachon (1997) allows a single supplier to choose operating policies for herself as well as all her retailers. It is assumed that the supplier chooses policies to optimize her own preferences, while leaving the retailers no worse off than they would be in the competitive solution. It is found that the supplier does not always choose the optimal policies. Nevertheless, the supplier frequently chooses better policies for the supply chain than the competitive solution. Hence, shifting control from one player to another does not eliminate all incentive conflicts, but often can mitigate them.

Narayanan and Raman (1997) study VMI between a supplier and a retailer. In addition to the supplier's product, the retailer sells a close substitute from another supplier. When a customer's favorite product is unavailable, the customer may switch to the other product which is always in stock. With VMI the retailer allows the supplier to choose his stock level of the supplier's product

in exchange for a fixed transfer payment. They demonstrate that VMI does not achieve the supply chain optimal profit. The problem is that the supplier stocks too much of her product because she does not consider the revenue the retailer earns when customers switch to the other product.

1.7 CONCLUSION

The literature on competitive supply chain inventory management recognizes that supply chains are usually operated by independent agents with individual preferences. Game theory is the methodological tool to determine how the players will behave when they each seek to maximize their own welfare. The key issues include whether a Nash equilibrium exists, whether there is a unique Nash equilibrium and whether the optimal policies ever belong to the set of Nash equilibria. It is frequently found that competitive and optimal behavior do not coincide, in which case it is worthwhile to investigate coordination techniques.

Some coordination techniques are designed to manipulate the behavior of one firm to the advantage of another. For example, buy-back and quantity discount contracts can be used by a supplier to increase her profits at the expense of the retailer's profits. Other techniques are designed to make the optimal policies incentive compatible, regardless of whether all of the players actually prefer the optimal solution over the competitive solution. Most of the techniques implement transfer payments between the players, but there is considerable variation in the form of these payments (e.g. linear fees and subsidies, two-part tariffs). It may also be possible to coordinate a supply chain by imposing service constraints on the parties, by shifting control among the players (vendor managed inventories), or by operating internal markets.

References

- Anderson, J., S. Axsater and J. Marklund, "Decentralized Multi-Echelon Inventory Control," Department of Industrial Engineering, Lund University, (1996).
- Anupindi, R. and Y. Bassok, "Centralization of Stocks: Retailers vs. Manufacturer," Northwestern University working paper, (1998).
- Axsater, S., "Exact and approximate evaluation of batch-ordering policies for two-level inventory systems," *Operations Research*, 41, 4 (1993), 777-785.
- Barnes-Schuster, D., Y. Bassok and R. Anupindi, "Coordination and Flexibility in Supply Contracts with Options," University of Chicago working paper, (1998).
- Boyaci, T. and G. Gallego, "Coordination Issues in Simple Supply Chains," Columbia University working paper, (1997).
- Butz, D., "Vertical Price Controls with Uncertain Demand," *Journal of Law and Economics*, 40, (1997), 433-459.
- Cachon, G., "Stock Wars: Inventory Competition in a Two Echelon Supply Chain," Fuqua School of Business working paper, (1997).
- Cachon, G. and M. Fisher, "Campbell Soup's Continuous Product Replenishment Program: Evaluation and Enhanced Decision Rules," *Production and Operations Management*, 6, 3 (1997), 266-276.
- Cachon, G. and M. Lariviere, "Capacity Choice and Allocation: Strategic Behavior and Supply Chain Performance," Fuqua School of Business working paper, Duke University, (1996).
- Cachon, G. and M. Lariviere, "Capacity Allocation with Past Sales: When to Turn-and-Earn," Duke University working paper, (1997a)
- Cachon, G. and M. Lariviere, "Contracting to Assure Supply, or What did the Supplier Know and When did He Know It?," Duke University working paper, (1997b).
- Cachon, G. and P. Zipkin, "Competitive and Cooperative Inventory Policies in a Two-Stage Supply Chain," Duke University working paper, (1997).
- Chen, F., "Decentralized Supply Chains Subject to Information Delays," Columbia University working paper, (1997).

- Chen, F., A. Federgruen and Y. Zheng, "Coordination Mechanisms for Decentralized Distribution Systems," Columbia School of Business working paper, (1997).
- Chen, F. and Y.-S. Zheng, "Lower Bounds for Multi-Echelon Stochastic Inventory Systems," *Management Science*, 40, 11 (1994), 1426-1443.
- Clark, A.J. and H.E. Scarf, "Optimal Policies for a Multi-Echelon Inventory Problem," *Management Science*, 6, (1960), 475-490.
- Clark, T. and J. Hammond, "Reengineering Channel Reordering Processes to Improve Total Supply Chain Performance," *Production and Operations Management*, 6, 3 (1997), 248-265.
- Corbett, C., "Stochastic Inventory Systems in a Supply Chain with Asymmetric Information," University of California at Los Angeles working paper, (1998).
- Deneckere, R., H. Marvel and J. Peck, "Demand Uncertainty, Inventories, and Resale Price Maintenance," *Quarterly Journal of Economics*, 111, (1996), 885-913.
- Deneckere, R., H. Marvel and J. Peck, "Demand Uncertainty and Price Maintenance: Markdowns as Destructive Competition," *American Economic Review*, 87, 4 (1997), 619-641.
- Donohue, K., "Supply Contracts for Fashion Goods: Optimizing Channel Profits," Wharton School working paper, (1996).
- Emmons, H. and S. Gilbert, "Returns Policies in Pricing and Inventory Decisions for Catalogue Goods," *Management Science*, 44, 2 (1998), 276-283.
- Eppen, G. and A. Iyer, "Backup Agreements in Fashion Buying - The Value of Upstream Flexibility," *Management Science*, 43, 11 (1997), 1469-1484.
- Federgruen, A. and P. Zipkin, "Computational Issues in an Infinite-Horizon, Multiechelon Inventory Model," *Operations Research*, 32, 4 (1984), 818-836.
- Fisher, M. and A. Raman, "Reducing the Cost of Demand Uncertainty through Accurate Response to Early Sales," *Operations Research*, 44, (1996), 87-99.
- Fudenberg, D. and J. Tirole (1991). *Game Theory*. Cambridge: MIT Press.
- Ha, A., "Supply Contract for a Short-Life-Cycle Product with Demand Uncertainty and Asymmetric Cost Information," Yale School of Management working paper, (1996).
- Hausman, W. and N. Erkip, "Multi-echelon vs Single-echelon Inventory Control Policies," *Management Science*, 40, 5 (1994), 507-602.
- Jeuland, A. and S. Shugan, "Managing Channel Profits," *Marketing Science*, 2, (1983), 239-272.
- Kandel, E., "The Right to Return," *Journal of Law and Economics*, 39, (1996), 329-56.
- Kouvelis, P. and M. Lariviere, "Decentralizing Cross-Functional Decisions: Coordination Through Internal Markets," Duke University working paper, (1997).
- Lal, R. and R. Staelin, "An Approach for Developing an Optimal Discount Pricing Policy," *Management Science*, 30, 12 (1984), 1524-1539.
- Lariviere, M. "Supply Chain Contracting and Coordination with Stochastic Demand," Fuqua School of Business working paper, (1998).

- Lee, H., P. Padmanabhan and S. Whang, "Information Distortion in a Supply Chain: The Bullwhip Effect," *Management Science*, 43, 4 (1997), 546-558.
- Lee, H. and M.J. Rosenblatt, "A Generalized Quantity Discount Pricing Model to Increase Supplier's Profits," *Management Science*, 32, 9 (1986), 1177-1185.
- Lee, H. and S. Whang, "Decentralized Multi-Echelon Supply Chains: Incentives and Information," Stanford University working paper, (1996).
- Mallik, S. and P. Harker, "Coordinating Supply Chains with Competition: Capacity Allocation in Semiconductor Manufacturing," University of Pennsylvania working paper, (1997).
- Milgrom, P. and J. Roberts, "Rationalizability, Learning and Equilibrium in Games with Strategic Complementarities," *Econometrica*, 58, (1990), 1255-78.
- Moorthy, K.S., "Managing Channel Profits: Comment," *Marketing Science*, 6, (1987), 375-379.
- Muckstadt, J.A. and L.J. Thomas, "Are Multi-Echelon Inventory Methods Worth Implementing in Systems with Low-Demand Rates?," *Management Science*, 26, 5 (1980), 483-494.
- Narayanan, V. and A. Raman, "Contracting for Inventory in a Distribution Channel with Stochastic Demand and Substitute Products," Harvard University working paper, (1997).
- Padmanabhan, V. and I.P.L. Png, "Returns Policies: Make Money by Making Good," *Sloan Management Review*, Fall, (1995), 65-72.
- Padmanabhan, V. and I.P.L. Png, "Manufacturer's Returns Policy and Retail Competition," *Marketing Science*, 16, 1 (1997), 81-94.
- Pasternack, B., "Optimal Pricing and Returns Policies for Perishable Commodities," *Marketing Science*, 4, 2 (1985), 166-176.
- Porteus, E., "Responsibility Tokens in Supply Chain Management," Stanford University working paper, (1997).
- Porteus, E. and S. Whang, "On Manufacturing/Marketing Incentives," *Management Science*, 37, 9 (1991), 1166-1181.
- Roundy, R., "98%-Effective Integer Ratio Lot-Sizing For One-Warehouse Multi-Retailer Systems," *Management Science*, 31, (1985), 1416-1430.
- Saggi, K. and N. Vettas, "On the Strategic Role of Fees and Royalties in Vertical Contracting," Duke University working paper, (1998).
- Spengler, J., "Vertical Integration and Antitrust Policy," *Journal of Political Economy*, (1950), 347-352.
- Sterman, J., "Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiment," *Management Science*, 35, 3 (1989), 321-339.
- Tirole, J. (1990). *The Theory of Industrial Organization*. Cambridge: MIT Press.
- Tsay, A., "The Quantity Flexibility Contract and Supplier-Customer Incentives," Santa Clara working paper, (1997).
- Tsay, A. and W. Lovejoy, "Supply Chain Control with Quantity Flexibility," Santa Clara University working paper, (1995).

Weng, K., "Channel Coordination and Quantity Discounts," *Management Science*, 41, (1995), 1509-1522.




Figure 1: Reaction functions, $\alpha = 0.3$, $p = 5$, $h_1 = h_2 = 0.5$, $L_1 = L_2 = 1$




Figure 2: Reaction functions, $\alpha = 0.9$, $p = 5$, $h_1 = h_2 = 0.5$, $L_1 = L_2 = 1$