STOCHASTIC INVENTORY SYSTEMS IN A SUPPLY CHAIN WITH ASYMMETRIC INFORMATION: CYCLE STOCKS, SAFETY STOCKS, AND CONSIGNMENT STOCK

CHARLES J. CORBETT

The Anderson School at UCLA, Box 951481, Los Angeles, California 90095-1481, charles.corbett@anderson.ucla.edu

(Received January 1998; revisions received November 1998, November 1999; accepted March 2000)

The two critical factors distinguishing inventory management in a multifirm supply-chain context from the more traditional centrally planned perspective are *incentive conflicts* and *information asymmetries.* We study the well-known order quantity/reorder point (Q, r) model in a two-player context, using a framework inspired by observations during a case study. We show how traditional allocations of decision rights to supplier and buyer lead to inefficient outcomes, and we use principal-agent models to study the effects of information asymmetries about setup cost and backorder cost, respectively.

We analyze two "opposite" models of contracting on inventory policies. First, we derive the buyer's optimal menu of contracts when the supplier has private information about setup cost, and we show how consignment stock can help reduce the impact of this information asymmetry. Next, we study consignment and assume the supplier cannot observe the buyer's backorder cost. We derive the supplier's optimal menu of contracts on consigned stock level and show that in this case, the supplier effectively has to *overcompensate* the buyer for the cost of each stockout.

Our theoretical analysis and the case study suggest that consignment stock helps reduce *cycle stock* by providing the supplier with an additional incentive to decrease batch size, but simultaneously gives the buyer an incentive to increase *safety stock* by exaggerating backorder costs. This framework immediately points to practical recommendations on how supply-chain incentives should be realigned to overcome existing information asymmetries.

1. INTRODUCTION

Traditionally, the multiechelon inventory control literature has assumed the existence of a central planner for the entire supply chain, with perfect information about cost structure and demand patterns and with the power to impose a globally optimal inventory policy on the echelons. More recently some authors have started dropping this assumption, attempting to explicitly capture the existence of multiple decision makers. A multifirm supply chain without a central planner differs from a centrally planned one in two fundamental ways. There will be *incentive conflicts* because different parties in the supply chain generally have different and often conflicting objectives. There will also be *information asymmetries* because individual parties have private information about, for example, cost structures or demand patterns, and no party has full information about the entire supply chain. Some existing work incorporates each of these individually, but this paper is one of the first to formally study the effect of both aspects simultaneously in an inventory control context. This leads to substantially different insights: Generally, when one considers only incentive conflicts or only information asymmetries, it is still feasible and optimal for one party to induce the other to choose the jointly optimal inventory policy through an appropriate incentive scheme. However, the presence of both issues combined destroys this property, leading to inefficient supply-chain performance.

It is important to realize that achieving joint optimality in itself need not be difficult. In either case, a principal aiming for jointly optimal outcomes could make the other party (who has full information) bear all costs, in exchange for a fixed payment. However, this is *not* optimal for the principal: The fixed payment he will have to offer to induce the agent to accept the contract depends on the agent's cost structure. Because this is unknown to the principal, the payment must be large enough to satisfy all possible types of agents, rendering this approach highly unattractive to the principal. (See Corbett and de Groote 2000 for more on this.) Therefore, any mechanism that optimizes the entire supply chain is either suboptimal for the principal or requires the principal to have full information. It is critical to bear this in mind because this is what sets the current framework apart from most previous work reviewed in §2.

Throughout this paper, we consider the widely used order quantity/reorder point system, under which an order of size Q is placed when the inventory position drops to the reorder point r. The classic model for this system, which is often referred to as the (Q, r) model, assumes that the demand is stochastic and stationary and unmet demand is backlogged. It also assumes that the system incurs a setup cost K for each order placement and a linear inventory holding cost h and backorder cost p per unit per period. The classic model assumes there is only one decision maker who bears all the costs and decides both Q and r. We transform this classic model into a *two-echelon* system with two firms,

Subject classifications: Inventory/production: asymmetric information supply chains. Game Theory: inventory control under asymmetric information. Inventory: lot-sizing and stochastic demand.
Area of review: Logistics & Supply Chain Operations.

a supplier and a buyer, involved in the decision-making process. We motivate our models with a recent case study in §3. To set benchmarks, we consider two extreme cases: the centralized system (the classic single decision-maker system) and the decentralized system (with no coordination at all). We discuss the centralized case in §4, where Zheng's (1992) analysis of the classic system is reviewed, and the decentralized case in §5, where the resulting behavior is highly inefficient. In §§6.1 and 7.1 we look at how either party could reduce this inefficiency by contracting on an inventory policy under full information.

To consider the situation where one party has private information about a certain parameter, we adapt the principal-agent contracting approach designed specifically for this type of adverse selection problems from Laffont and Tirole (1993). We study two different situations. First, in §6, we analyze the case where the buyer is the principal and the supplier has private information about his setup costs K, and we derive the buyer's optimal menu of contracts. We also analyze special cases, one of which is the classical EOQ formula. Then, in §7, we focus on the case with the supplier as principal, but where the buyer has private information about backorder costs p, and we derive the supplier's optimal menu of contracts. A special case here is the classical newsboy formula. We also examine whether contracting on a stockout penalty, to be paid by the supplier to the buyer, instead of contracting directly on an inventory policy, would lead to more efficient results. We find that it is equivalent to contracting directly on (Q, r); however, the supplier needs to *overcompensate* the buyer for the cost of individual stockouts. The cases and results are summarized in Table 2. Finally, in §8 we summarize our conclusions and offer suggestions for further research.

The findings in this paper lead to several critical insights into supply-chain performance. It demonstrates the importance of explicitly considering incentive and information issues in modeling supply chains. Together, they generally lead to inefficient outcomes, even under an optimal menu of contracts. This paper shows how economic models of incentive conflicts and information asymmetries can be formally superimposed upon traditional inventory models while still allowing tractable analysis. It suggests that, in (re-)designing supply chains, one should consider the incentive effects and information requirements and attempt to agree on an incentive structure that overcomes these inefficiencies. Specifically, we show how, in the coordination framework followed here, consignment stock is a simple and robust mechanism that can help reduce the inefficiency caused by the supplier choosing a lot size Q that is greater than the jointly optimal lot size, even when the buyer cannot observe his setup cost K. However, consignment stock leads to a higher reorder point r than is jointly optimal when the buyer has private information about backorder cost p. In other words, in this context consignment stock can *reduce* the impact of information asymmetries on cycle stocks but can *exacerbate* their impact on safety stocks.

In practical terms, this means that in supply chains where long production cycles are the leading cause of inventory, the supplier should be made to bear (at least part of) the associated cost, for example, through a consignment scheme. Conversely, if downstream demand uncertainty is the key cause of inventory, the buyer should bear the associated cost of safety stock. In other words, the decision rights throughout the supply chain should be in alignment with the costs associated with those decisions and with the information needed to make those decisions appropriately. This paper shows how misalignments can easily occur in practice, explores the impact of such misalignments, and proposes incentive contracts to help overcome the resulting inefficiencies.

2. LITERATURE

The classic work on designing, rather than operating, an inventory policy using only locally available information is Clark and Scarf (1960): They derive a decomposition scheme for a multiechelon stochastic inventory system, which, as Axsäter and Rosling (1993) point out, can be implemented using only local inventory information. Lee and Whang (1996) show that Clark and Scarf's decomposition approach can be implemented in an "informationally decentralized" fashion by a combination of consignment, delivery warranties, and an additional backlog penalty at the downstream site. Chen (1999) adds information delays to the same multiechelon framework and designs incentive schemes based on payments between echelon managers and the firm rather than directly between echelon managers, as in Lee and Whang (1996). All three do, however, require the presence of a central planner to determine and enforce the decomposition scheme. Our basic model is somewhat simpler, but we show how, in the absence of a central planner, contracting can no longer achieve joint optimality.

It is well understood that different parties in a supply chain face different incentives, and that, left to their own devices, they will generally not arrive at a system-wide optimum. An extensive review is given by Tsay et al. (1999). In Cachon and Larivière (1997), a manufacturer faces uncertain demand and contracts with a supplier who builds capacity in anticipation of demand realization. They study full and asymmetric information (where the supplier cannot observe the demand distribution) and forced and voluntary compliance (referring to whether the supplier is obliged to actually ship the quantity agreed on). They find that the value of each type of contract varies with the information structure and compliance regime in place. Cachon (1997) compares three coordination mechanisms for reorder points in a one-supplier, multiretailer context. There, the jointly optimal stocking policies are often not a Nash equilibrium, and the only way to enforce them is to change the player's incentive structure; however, selection of the best among several equilibria or centralizing control at the supplier (through vendor-managed inventory combined with side payments) can lead to improvements.

In the one-supplier, one-retailer case, Cachon and Zipkin (1999) let buyer and supplier share backorder costs according to a ratio α; they find that the unique Nash equilibrium is (essentially) never the optimal solution, and that the joint efficiency loss increases with the difference in backorder penalty incurred by the two parties. In particular, if the supplier fully compensates the buyer for his backorder costs, coordination will not occur; our findings in §7.5 are analogous.

The basic premise of the joint economic lot sizing (JELS) literature—including Monahan (1984), Lal and Staelin (1984), Dada and Srikanth (1987), and Weng (1995)—is that, in an EOQ-type environment, both buyer and supplier can be better off by coordinating on the joint optimum. A critical and questionable assumption made throughout that literature, though, is that all parties have full information about cost parameters. Corbett and de Groote (1997, 2000) show how the equivalence of a range of common coordination schemes breaks down under asymmetric information. Here, we push the coordination principle of the JELS literature a step further by explicitly modeling such information asymmetry.

In a paper closely related to ours, Ha (1997) studies a supplier with private information about his manufacturing cost and a customer facing stochastic demand. Ha determines the customer's optimal order policy and shows that under asymmetric information it is no longer optimal for the customer to induce the system-wide optimal solution. We do not consider price-sensitive demand here but do include setup costs, and we do allow both supplier and buyer to hold private information. Cachon and Larivière (1999) ask how a supplier should allocate scarce inventory to multiple retailers when the supplier cannot observe the retailers' inventory position. They compare various allocation mechanisms, some of which will lead retailers to inflate their orders, and show that, surprisingly, system-wide performance need not be hurt by such manipulation because the supplier may be induced to choose a higher capacity level than under a truth-inducing allocation mechanism.

To capture some of the dynamics crucial to understanding buyer-supplier relationships, we draw heavily on Laffont and Tirole (1993), especially their principal-agent models of adverse selection.

3. CASE EXAMPLE: INTRODUCING CONSIGNMENT STOCK

Let us start by briefly describing the case that inspired the models in this paper. (The context is described in more detail in Corbett 1996 and Corbett et al. 1999; the description given here is simplified somewhat to avoid unnecessary complexity.) The supplier is a large multinational, selling chemical products to automotive suppliers. The supplier was often sluggish in responding to customers' demands, as a result of the long production cycles often found in heavy process industries. This was a frequent cause for customer complaints. The supplier took the initiative to transform their traditionally rather distant relationship with a key customer toward a partnership-type relationship. Both firms represented a substantial proportion of each other's volume, but neither dominated the relationship. A commercial "partnership agreement" was reached, specifying (among other terms) a price concession in exchange for a volume commitment: The customer agreed to allocate a given annual volume of their business to the supplier. This volume commitment allows us to ignore competition in modeling at least in the short or medium term. The agreement also included consignment stock: The supplier would own the inventory held at the customer's site until it was consumed. The customer would pull whatever material they needed; after each production cycle, the supplier would replenish up to a base level previously agreed upon. The base level would have to be enough to cover anticipated demand until completion and delivery of the next production lot and to provide adequate safety stock. As a first effect of the consignment agreement, the supplier became anxious about the additional inventory burden they would have to carry; an internal e-mail message stated that "to make the whole thing work, cycle time reduction in our plant will be critical."

The buyer, as is typical of automotive suppliers, faces a huge stockout cost and therefore held enough incoming material stock to cover at least six weeks of demand (the supplier's production cycle time plus two weeks' worth of safety stock). The supplier expected that because of the volume commitment and cycle time reduction they could adopt an order-up-to level, which would be much lower than the buyer's current six weeks' worth of stock. Total inventory costs (of supplier and buyer combined) could then be reduced. (The buyer no longer cares about holding costs as inventory is now held on consignment by the supplier.) In theory, the supplier would determine the jointly optimal order-up-to level, taking into account all inventory costs (holding costs and stockout costs) for both buyer and supplier. The supplier could not determine precisely what this jointly optimal order-up-to level should be, however, because the buyer's precise stockout cost (whether explicit in the form of a penalty or implicit through a negative mark on their vendor rating) is hard to measure and certainly is not known to the supplier. In other words, whatever order-up-to level or fill rate the supplier proposed, the buyer could argue it should be higher in light of this enormous stockout cost. In practice, the implied stockout cost varies from order to order: Most orders were routine and would not cause severe disruption if delivered slightly late, whereas others were designated "rush orders" by the buyer, implying that nondelivery would truly invoke the full stockout cost. The supplier, however, cannot verify whether a so-called rush order is truly urgent or merely a precautionary measure allowing the buyer to cushion the effects of poor internal production and purchasing discipline. In several situations orders designated as "urgent" (implying a high stockout

Table 1. Summary of notation.

r, Q	Reorder point and order quantity
$r^{\max}, Q^{\max}$	Upper bound on reorder point and order quantity (used only in §5)
V	Stockout penalty paid by supplier to buyer (§7.5)
X_j, X_s, X_b	Jointly optimal, supplier's preferred, and buyer's preferred value for any variable X
L	Delivery lead time
$\Psi(z), \mu, \sigma$	Distribution of lead time demand z, with expectation μ and standard deviation σ
$[\underline{z}, \bar{z}]$	Domain (not necessarily finite) of demand distribution
p	Buyer's backorder cost per unit per unit time
h	Holding cost per unit per unit time
K	Supplier setup cost
C_j, C_s, C_b	Joint, supplier, and buyer cost function
$G(y)$	Rate at which costs accumulate at time $t+L$ if current inventory position is y
$G_h^+(y), G_p^-(y)$	Rate at which holding costs and backorder costs, respectively, accumulate at time $t+L$, when unit holding cost rate is h and unit backorder cost rate is p and inventory position at time t is y
P	Per-period compensation payment in menu of contracts
$C_s^{\max}, C_b^{\max}$	Maximum acceptable supplier and buyer costs after compensation
$F_i(\cdot), f_i(\cdot)$	Prior distribution and density over parameter i
$\dot{X}(y) = D_y X(y)$	Derivative of $X(y)$ with respect to y
$D_i X(y, z)$	ith partial derivative of $X(y, z)$
$D_{ij} X(y, z)$	Second-order derivative of $X(y, z)$, i.e., $D_{ij}X(y, z) = D_i D_j X(y, z)$
$E_F[X]$	Expectation of X with respect to distribution F
$(x)^+$	$\max\{0, x, \}$
$[\underline{x}, \bar{x}]$	Lower and upper bound of any variable x
$\hat{x}(x)$	Announced value of a parameter when true value is x

cost) were later found not to be urgent but merely precautionary in nature. Through this and other mechanisms, the buyer resisted the supplier's attempts to reduce the level of consigned stock.

The case illustrates the challenge of determining an appropriate inventory level, when the inventory in question is simultaneously a cycle stock—caused by the supplier's cycle time—and safety stock, to cover the demand uncertainty faced by the buyer. Below we model a similar context where two players both have some influence over total inventory cost. We show that the total costs and those of both players individually depend on the allocation of decision rights and costs among the players, just as they did in the case.

4. THE BASIC (Q, r) MODEL IN A CENTRALIZED SETTING

4.1. Assumptions and Notation

Throughout this paper we use the (Q, r) lot size/reorder point model, adapted to a two-player context, as follows. The supplier sets the production lot size Q, while the buyer determines the reorder point r. The single-player version is analyzed in detail in Zheng (1992) and Zheng and de Groote (1993). The key qualitative insights on the incentive effects of consignment stock are not sensitive to the precise setting chosen. We make the following assumptions (notation largely follows that in Zheng 1992 and is summarized in Table 1). There is one buyer, buying a single product from one supplier, who delivers in batches of size Q, with a constant delivery lead time L. The supplier's setup costs are K per batch. Transfer lots are equal to production lots, representing, for instance, a make-to-order environment. The buyer is faced with stochastic demands, arriving at rate λ. Supplier delivery lead time is L, and lead time demand z follows a distribution $\Psi(\cdot)$ on $[0, \infty[$. We follow Zheng (1992) in assuming cumulative demand can be modeled by a nondecreasing stochastic process with stationary increments and continuous sample paths so that inventory position in steady state will be uniformly distributed on $]r, r+Q]$ and independent of lead time demand z. (Zheng refers to Serfozo and Stidham 1978 and Browne and Zipkin 1989 for a detailed discussion of this assumption.) All payments take place immediately upon delivery of the goods; the only holding costs are the opportunity costs of capital. All stockouts are backordered. Holding costs accrue to either supplier or buyer at a constant rate h per unit stock per unit time; backorder costs, borne by the buyer, accrue at a constant rate p per unit backorder per unit time. Many of these assumptions can

Table 2. Overview of results.

	Section 6	Section 7	Section 7.5
"Principal," i.e., party with initiative	Buyer	Supplier	
"Agent," i.e., party with private information	Supplier	Buyer	
Information asymmetry about	Setup cost K	Backorder cost p	
Initial stock ownership	No consignment stock	Consignment stock	
Contracting on	Inventory policy (Q_1, r_1, P_1)	Inventory policy (Q_2, r_2, P_2)	Supplier stockout penalty (V_3, P_3)
Finding	An incentive contract that includes consignment stock *reduces* inefficiency caused by information asymmetry.	Consignment *exacerbates* inefficiency caused by information asymmetry; an incentive contract that passes part of the holding costs back to the buyer reduces this inefficiency.	Consignment *exacerbates* inefficiency caused by information asymmetry; contracting can help, but supplier must *overcompensate* buyer for stockouts.
Special case	Deterministic demand: EOQ-type expression	No setup costs: newsboy-type expression	

easily be relaxed; one can include fixed ordering costs for the buyer or physical holding costs for both parties, but this does not enhance the insights.

4.2. Joint Optimum

We follow Zheng's (1992) formulation of the (Q, r) model. Let $G(y)$ be the rate at which expected holding and backorder costs accumulate at time $t+L$ when the inventory position at time t equals y. Then,

$$G(y) = E_\Psi[h(y-z)^+ + p(z-y)^+]. \tag{1}$$

Later we will need to distinguish the holding cost and backorder cost components:

$$G_h^+(y) = E_\Psi[h(y-z)^+], \tag{2}$$

$$G_p^-(y) = E_\Psi[p(z-y)^+]. \tag{3}$$

The following properties of $G_p^-(\cdot)$ are easy to verify and are used in the proofs in the appendix:

$$\alpha G_\beta^-(y) = G_{\alpha\beta}^-(y), \quad \forall\, \alpha, \beta, y, \tag{4}$$

$$D_\beta G_\beta^-(y) = G_1^-(y), \quad \forall\, \beta, y, \tag{5}$$

where D_β denotes the derivative with respect to β. Zheng (1992) bases his analysis on the average joint cost rate $C_j(Q, r)$, in the following form:

$$C_j(Q, r) = \frac{\lambda K + \int_r^{r+Q} G(y)dy}{Q}. \tag{6}$$

The first term is the rate at which setup costs accrue, the second is the expected holding and backorder cost rate, where the expectation is simply the average holding and backorder cost rate over the interval $]r, r+Q]$ due to Zheng's (1992) assumption on the demand process.

Let r_j and Q_j denote joint optima. Zheng (1992) uses a sequential approach, by first optimizing r for given Q, and then analyzing the resulting cost function $C_j(Q) := C_j(Q, r_j(Q))$. Write $D_Q r(Q)$ for the derivative; the following lemma then summarizes some results (largely from Zheng 1992 and Zheng and de Groote 1993) for $r_j(Q)$ and Q_j and defines cycle stock and safety stock, to which we return in the special cases studied later in this paper.

LEMMA 1. (a) *For any $Q > 0$, $r = r_j(Q)$ if and only if $G(r) = G(r+Q)$.*

(b) *For any $Q > 0$, $r_j(Q)$ is solely determined by $G(\cdot)$ and hence independent of K.*

(c) *For any $Q > 0$, $-1 < D_Q r_j(Q) < 0$, so $r_j(Q)$ is decreasing in Q.*

(d) *$r_j(Q_j)$ is increasing in p.*

(e) *Q_j is increasing and r_j decreasing in K.*

(f) $\lim_{Q\uparrow\infty} r(Q) = -\infty$.

(g) *If demand $\Psi(\cdot)$ is deterministic, then*

$$Q_j = \sqrt{\frac{2\lambda K(h+p)}{hp}}, \tag{7}$$

$$r_j = \lambda L - \frac{h}{h+p} Q_j. \tag{8}$$

The reorder point r_j may be negative as it is based on inventory position (on-hand plus on-order minus backorders). All system inventory is now induced by the production lot size Q, and is therefore cycle stock. *The average cycle stock is given by $\frac{1}{2}\frac{p-h}{p+h}Q_j$. (Note that if $p < h$, the system will, on average, have backorders.)*

(h) *If setup costs $K = 0$, then $Q_j = 0$ and $r_j = \Psi^{-1}(\frac{p}{h+p})$. All inventory in the system is now induced by both the demand uncertainty and the delivery lead time L that are inherent in Ψ. The excess inventory above λL (the amount held exclusively to cover delivery lead time L) is the* safety stock. *So, safety stock here is equal to $\Psi^{-1}(\frac{p}{h+p}) - \lambda L$. (Note that this may be negative, depending on p, h, and Ψ.)*

5. THE DECENTRALIZED (Q, r) MODEL WITHOUT CONTRACTING

5.1. Definition of the Decentralized Setting

In this section we formally introduce the notion that the (Q, r) system described above is being managed by two players. Initially we do not let them coordinate in any way, from §6 onward we introduce the possibility of contracting. First, we formally define the games, Γ when there is no consignment stock, Γ^c when there is consignment stock. The players are the supplier and the buyer; their payoff functions depend on whether or not we assume consignment stock. In game Γ without consignment stock, the respective cost functions are

$$C_s(Q, r) = \frac{\lambda K}{Q}, \quad C_b(Q, r) = \frac{\int_r^{r+Q} G(y)dy}{Q}. \tag{9}$$

In game Γ^c with consignment stock, the respective cost functions are

$$\begin{aligned} C_s^c(Q, r) &= \frac{\lambda K + \int_r^{r+Q} G_h^+(y)dy}{Q}, \\ C_b^c(Q, r) &= \frac{\int_r^{r+Q} G_p^-(y)dy}{Q}. \end{aligned} \tag{10}$$

Note that we do not need the lot-for-lot assumption under consignment. The supplier's strategy is Q, the buyer's is r, with $(Q, r) \in [0, Q^{\max}] \times [-Q^{\max}, r^{\max}]$. We assume $Q^{\max}$ and r^{max} are sufficiently large not to be restrictive in §6 and beyond. As in Cachon (1997), this lower bound on r is not restrictive.

It is easy to see that without coordination of some sort, both games are poorly behaved. With no consignment, the supplier will want to produce and ship the largest lot size possible. Similarly, under consignment, the buyer will wish to have the highest possible level of inventory. In both cases, the equilibria are not constrained by anything other than the artificial bounds $Q^{\max}$ and $r^{\max}$, making these games ill-behaved. One could ask whether a different allocation of decision rights, e.g., letting the supplier or the buyer set both Q and r, with or without consignment, would lead to more desirable outcomes. It is easy to verify that this is not the case. In the no-consignment game Γ, the key problem is that the supplier would not internalize the holding costs associated with a large reorder quantity Q, whereas the buyer would not internalize the setup costs associated with a small Q. In the consignment game Γ^c, the buyer would not internalize the holding costs associated with a high reorder point r, whereas the supplier would not internalize the stockout costs associated with a low r. Clearly, if one goes another step further and gives one party the right to set both Q and r and also makes that party bear all setup, holding, and stockout costs, that party will choose the jointly optimal solution. Below we will see that under asymmetric information this incentive scheme is not optimal for either party.

In the absence of some form of coordination the resulting outcomes will be highly inefficient. The contracting approach reduces these inefficiencies by fundamentally changing how decision rights are allocated: One party now establishes a rule linking Q and r (the "menu of contracts"); the other party, holding private information, is then given the right to choose any (Q, r) combination from that menu. This way, both parties still take part in the decision-making process, but Q and r are no longer determined independently of one another.

6. CONTRACTING ON Q UNDER INFORMATION ASYMMETRY ABOUT K

6.1. Coordinating on Lot Size Q Under Full Information

In this section we start with the situation without consignment stock. The buyer will already choose $r_b = r_j(Q)$ for whatever Q the supplier chooses. It is clear that by coordinating on the jointly optimal lot size Q_j rather than letting the supplier choose some far higher Q_s, as he would do without coordination, both supplier and buyer could be made better off. Analogously to the basic principle of the joint economic lot sizing literature (see, e.g., Monahan 1984 or Lal and Staelin 1984), we know that in the no-consignment game Γ, both parties can be made better off by coordinating on $(Q_j, r_b(Q_j))$; i.e., $C_b(Q, r_j(Q)) - C_b(Q_j, r_j(Q_j)) \geqslant C_s(Q_j, r_j(Q_j)) - C_s(Q, r_j(Q)) \forall Q$. The buyer could implement this by offering the supplier a contract $(Q_j, r_j(Q), P)$ specifying a compensation rate $P \in [C_s(Q_j, r_j(Q_j)) - C_s(Q, r_j(Q)), C_b(Q, r_j(Q)) - C_b(Q_j, r_j(Q_j))]$ in return for the supplier's choosing $Q = Q_j$. Under full information, the buyer would set the compensation rate $P = C_s(Q_j, r_j(Q_j)) - C_s(Q, r_j(Q))$, keeping all remaining savings for himself. In fact, if the buyer can offer such a side payment P, he can do even better: Let $C_s^{\max}$ be the highest total net costs (after side payment) under which the supplier is still willing to contract with the buyer. With full information, the buyer can impose a side payment $P = C_s(Q_j, r_j(Q_j)) - C_s^{\max}$, pushing the supplier to incur the highest costs he will accept, and keeping all remaining cost savings for himself. However, such a compensation P depends on $C_s(\cdot)$, so the buyer must know the supplier's setup cost K. This is a strong assumption, and below we show how the buyer can offer a menu of contracts when he cannot observe K.

6.2. The Buyer's Optimization Problem Under Asymmetric Information

To address this adverse selection problem, we follow the approach outlined in Laffont and Tirole (1993). The buyer is the party taking the initiative and therefore is the principal, but the supplier holds private information about his setup cost K. Assume the buyer holds a prior distribution

$F_K(\cdot)$, differentiable on its domain $[\underline{K}, \overline{K}]$ with $0 < \underline{K} < \overline{K} < \infty$ with density $f_K(\cdot)$, and such that $F_K(\cdot)$ satisfies the monotone decreasing reversed hazard rate property:

Decreasing reversed hazard rate assumption:

$$D_K \frac{f_K(K)}{F_K(K)} \leqslant 0.$$

For more on this assumption, see Appendix B. Instead of proposing a single contract (Q, r, P) as above, the buyer now offers a *menu of contracts* $\{Q, r(Q), P(Q)\}$, letting the supplier choose a specific $(Q, r(Q), P(Q))$-triplet from the menu. To model the contracting process, we parameterize each of Q, r, and P on K. Note that offering a $\{Q(K), r(K), P(K)\}$ menu is equivalent to a $\{Q, r(Q), P(Q)\}$ menu, although that equivalence relation need not exist in closed form. Whether or not contracting is explicitly based on K, the supplier's setup cost K can be inferred from the $(Q, r(Q), P(Q))$ triplet he selects. The contracting steps are:

Step 1. The supplier learns his setup cost K, the buyer does not.

Step 2. The buyer offers menu $\{Q(\cdot), r(\cdot), P(\cdot)\}$, linking $Q(\widehat{K})$ and $r(\widehat{K})$ to compensation rate $P(\widehat{K})$ for whatever setup cost $\widehat{K}$ the supplier announces.

Step 3. The supplier chooses a contract $(Q(\widehat{K}), r(\widehat{K}), P(\widehat{K}))$, effectively announcing $\widehat{K}$.

Step 4. Production and delivery lot size are fixed (forever) at $Q(\widehat{K})$, reorder point at $r(\widehat{K})$, and the supplier receives compensation rate $P(\widehat{K})$. Buyer and supplier incur net costs at per-period rates $C_b(Q(\widehat{K}), r(\widehat{K})) + P(\widehat{K})$ and $C_s(Q(\widehat{K}), r(\widehat{K})) - P(\widehat{K})$, respectively.

We make two common assumptions. First, we assume that contracting is a one-shot process, so that renegotiation is not allowed. Second, once the supplier has announced his choice of $Q(\widehat{K})$ and $r(\widehat{K})$, the buyer can verify and enforce this inventory policy. This is clearly true here, as the buyer receives lots of size $Q(\widehat{K})$ and controls the reorder point $r(\widehat{K})$ himself.

From the *revelation principle* (see, e.g., Fudenberg and Tirole 1991, p. 256; Laffont and Tirole 1993, p. 120; or Myerson 1979) we can restrict our attention to menus of contracts in which the supplier announces his setup cost truthfully. Intuitively, the buyer can predict what setup cost $\widehat{K}$ any supplier with true cost K would announce, so he can take this knowledge of $\widehat{K}(K)$ into account in designing the menu $\{Q(\cdot), r(\cdot), P(\cdot)\}$. This allows us to formulate an incentive-compatibility constraint, which ensures that a supplier with setup cost K will indeed announce $\widehat{K} = K$. For derivatives with respect to the parameter being contracted on, write $\dot{Q}(K) = D_K Q(K)$, etc. The optimization problem $\mathscr{S}_1$ faced by a supplier with true setup cost K and having to choose what setup cost $\widehat{K}$ to announce is

$$\mathscr{S}_1 \quad \min_{\widehat{K}} \; \frac{\lambda K}{Q(\widehat{K})} - P(\widehat{K}). \tag{11}$$

The first-order conditions require $\frac{\lambda K}{Q(\widehat{K})^2}\dot{Q}(\widehat{K}) + \dot{P}(\widehat{K}) = 0$. Inducing truthful revelation is equivalent to ensuring that the optimum is reached at $\widehat{K} = K$, hence we can formulate

Incentive-compatibility constraint on $P(\cdot)$:

$$\textbf{IC:} \quad \dot{P}(K) = -\frac{\lambda K}{Q(K)^2}\dot{Q}(K) \quad \forall\, K \in [\underline{K}, \overline{K}]. \tag{12}$$

More precise and complete proofs, including verification of sufficiency conditions, for the preceding and following expressions are given in Appendix C (except for the model examined in §7, as that model is slightly more complex). To make the contract acceptable to the supplier, regardless of K, total costs (after compensation) under the contract may not exceed some predetermined and commonly known limit $C_s^{\max}$. This is formulated as follows:

Participation constraint:

$$\textbf{PC:}\; C_s^{\max} \geqslant C_s(Q(K), r(K)) - P(K) \quad \forall\, K \in [\underline{K}, \overline{K}]. \tag{13}$$

To find the optimal menu of contracts, the buyer solves the following problem:

$$\begin{aligned} \mathscr{B}_1 \quad \min_{Q(\cdot),\, r(\cdot),\, P(\cdot)} \quad & E_{F_K}[P(K) + C_b(Q(K), r(K))], \\ \text{s.t.} \quad & \text{IC (12), PC (13).} \end{aligned} \tag{14}$$

We are not interested in the actual value of $C_s^{\max}$, as this will only shift $P(\cdot)$ upward or downward; we focus on the shape of $P(\cdot)$. To see this, note that objective function (14) and constraint (12) jointly determine the optimal $\dot{P}(\cdot)$, so that $P(\cdot)$ itself is determined up to a constant of integration. This constant, in turn, is set using participation constraint (13).

6.3. The Buyer's Optimal Menu of Contracts

Solving the optimal control problem $\mathscr{B}_1$ yields the optimal menu $\{Q(\cdot), r(\cdot), P(\cdot)\}$ shown below.

PROPOSITION 1. *Assume that the first-order conditions for $\mathscr{B}_1$ are also sufficient. Then the optimal menu of contracts $\{Q_1(\cdot), r_1(\cdot), P_1(\cdot)\}$, when the supplier holds private information about K, is such that*:

$$Q_1(K) = Q_j\left(K + \frac{F_K(K)}{f_K(K)}\right) \quad \forall\, K \in [\underline{K}, \overline{K}], \tag{15}$$

$$r_1(K) = r_j(Q_1(K)) \quad \forall\, K \in [\underline{K}, \overline{K}], \tag{16}$$

$$\dot{P}_1(K) = -\frac{\lambda K}{Q_1(K)^2}\dot{Q}_1(K) \quad \forall\, K \in [\underline{K}, \overline{K}]. \tag{17}$$

$Q_1(K) \geqslant Q_j(K)$ and $Q_1(\cdot)$ is increasing in K, $r_1(\cdot)$ is decreasing in K, $P_1(\cdot)$ is decreasing in K. Moreover, the supplier's net costs after compensation, $C_s(Q_1(K), r_1(K)) - P_1(K)$, are increasing in K.

Proposition 1 shows that $(Q_1(\cdot), r_1(\cdot))$ is always less efficient than $(Q_j(\cdot), r_j(\cdot))$: $Q_1(K) \geqslant Q_j(K)$ and $C_j(Q_1(K), r_1(K)) \geqslant C_j(Q_j(K), r_j(K))$ for all $K \in [\underline{K}, \overline{K}]$. Moreover, at $K = \underline{K}$, i.e., the supplier with the lowest possible setup cost K, $F_K(\underline{K}) = 0$, so that $Q_1(\underline{K}) = Q_j(\underline{K})$. The inefficiency depends on $\frac{F_K(K)}{f_K(K)}$, which increases in K. After revealing K, both parties could still be better off coordinating on the joint optimum, but with renegotiation, the relevation principle and thus mathematical tractability are lost.

6.4. Deterministic Demand: Contracting to Reduce Cycle Stock

Assuming for a moment that demand $\Psi(\cdot)$ is deterministic, we can relate the above expressions to the EOQ formula. All inventory in the system is now induced by the production lot size Q, and is therefore cycle stock.

COROLLARY 1. *Consider $\mathcal{B}_1$ with demand $\Psi(\cdot)$ deterministic. Assume that the first-order conditions for $\mathcal{B}_1$ are sufficient. Then the optimal menu of contracts $\{Q_1^d(\cdot), r_1^d(\cdot), P_1^d(\cdot)\}$, when the supplier holds private information about K, is such that*:

$$Q_1^d(K) = \sqrt{2\lambda\left(K + \frac{F_K(K)}{f_K(K)}\right)\frac{h+p}{hp}} \quad \forall K \in [\underline{K}, \overline{K}], \quad (18)$$

$$r_1^d(K) = \lambda L - \frac{h}{h+p} Q_1^d(K) \qquad \forall K \in [\underline{K}, \overline{K}], \quad (19)$$

$$\dot{P}_1^d(K) = -\frac{1}{2}\frac{K}{K + \frac{F_K(K)}{f_K(K)}}\frac{hp}{h+p}\dot{Q}_1^d(K)$$
$$= -\frac{K}{K + \frac{F_K(K)}{f_K(K)}} \cdot D_K C_b(Q_1^d(K), r_1^d(K))$$
$$\forall K \in [\underline{K}, \overline{K}]. \quad (20)$$

In addition to the properties listed in Proposition 1, $P_1^d(\cdot)$ *is convex decreasing in* K.

This helps us to further interpret the results: At $K = \underline{K}$, we have $\dot{P}_1^d(\underline{K}) = -\frac{1}{2}\frac{hp}{h+p}\dot{Q}_1^d(\underline{K}) = -D_K C_b(Q_1^d(\underline{K}), r_1^d(\underline{K}))$ (taking right-hand derivatives). Were a supplier with any setup cost K to announce a marginally higher setup cost, the buyer's costs would increase by $D_K C_b(Q_1^d(K), r_1^d(K))$, so at $K = \underline{K}$ the buyer passes this increase on to the supplier *in full*. For $K > \underline{K}$, $0 > \dot{P}_1^d(K) > D_K C_b(Q_1^d(K), r_1^d(K))$, so the buyer passes a smaller portion of his cost increase back to the supplier. This corresponds to Laffont and Tirole's (1993) findings: The principal (the buyer) will present more efficient agents with more "high-powered" incentives, and move more toward a cost-plus reimbursement scheme for less efficient agents. In the extreme case $K = \underline{K}$, the incentive scheme is that of consignment combined with a delivery warranty (the supplier now incurs both holding costs and backorder costs).

6.5. Consignment Stock and Vendor-Managed Inventory (VMI)

To relate the above expressions directly to the simplest form of the EOQ model, assume backorders are not allowed (or, equivalently, let p go to ∞). Then we find

$$Q_1^d(K) = \sqrt{\frac{2\lambda\left(K + \frac{F_K(K)}{f_K(K)}\right)}{h}} \quad \forall K \in [\underline{K}, \overline{K}], \quad (21)$$

$$r_1^d(K) = \lambda L \quad \forall K \in [\underline{K}, \overline{K}], \quad (22)$$

$$\dot{P}_1^d(K) = -\frac{1}{2}h\frac{K}{K + \frac{F_K(K)}{f_K(K)}}\dot{Q}_1^d(K) \quad \forall K \in [\underline{K}, \overline{K}]. \quad (23)$$

Safety stock in this case is zero, as indeed it should be: We had defined safety stock as $\Psi^{-1}(\frac{p}{p+h}) - \lambda L$, and as lead time demand is known with certainty to be equal to λL, we have $\Psi^{-1}(\frac{p}{p+h}) = \lambda L$ for any $\frac{p}{p+h} > 0$. Now, at $K = \underline{K}$, $\dot{P}_1^d(\underline{K}) = -\frac{1}{2}h\dot{Q}_1^d(K)$, so the increased holding costs resulting from announcing a higher K are passed on to the supplier *in full*. This is equivalent to a consignment scheme: The buyer does not pay the supplier for the goods until they are sold to the final customer, making the supplier bear the financial holding costs for incoming material held at the buyer's site. For $K > \underline{K}$, the buyer passes a decreasing proportion of the marginal increase in holding costs back to the supplier. In other words, the more favorable the supplier's cost structure (the lower his setup costs K), the less he is "punished" by the buyer. Were the buyer to impose consignment for all K, the corresponding compensation payment would become $\dot{P}_c^d(K) = -\frac{1}{2}h\dot{Q}_c^d(K)$, so that $P_c^d(K) = -\frac{1}{2}hQ_c^d(K) + P_c^0$ for some constant P_c^0. It is easy to verify that the supplier would then choose $Q_c^d(K) = Q_j^d(K)$. In other words, under deterministic demand and when backorders are not allowed, consignment stock will induce the supplier to choose the jointly optimal reorder quantity despite the information asymmetry around K.

This is a good point to elaborate on the two reasons why the buyer's expected costs after contracting are higher than they would be under full information. System-wide costs are higher because the inventory policy (Q_1, r_1) resulting from contracting is not equal to the jointly optimal inventory policy (Q_j, r_j) that the buyer would choose to enforce under full information. The difference between total system costs in the two cases is the *inefficiency* caused by the information asymmetry. In addition, under asymmetric information the buyer has to bear a larger share of total system costs. Under full information, the buyer can force the supplier to incur his maximum acceptable cost level $C_s^{\max}$, using an appropriate mechanism with side payment; under asymmetric information, the buyer must accept that a supplier with setup costs lower than the worst-case $\overline{K}$ incurs net costs (after side payment) lower than $C_s^{\max}$. The net cost reduction (after side payment) available to the supplier due to the information asymmetry is known as the *information rent* he can extract.

The total cost increase to the buyer as a result of the information asymmetry is the sum of the system inefficiency and the supplier's information rent above. More formally: The system inefficiency is equal to $C_j(Q_1(K), r_1(K)) - C_j(Q_j(K), r_j(K))$. The supplier's information rent is the difference between his net total costs when the buyer has full information ($C_s^{\max}$) and those after contracting under asymmetric information, given by $C_s^{\max} - [C_s(Q_1(K), r_1(K)) - P(K)]$. The total cost of information asymmetry to the buyer is the difference between his total net costs after contracting under asymmetric information and those under full information and is $[C_b(Q_1(K), r_1(K)) + P(K)] - [C_j(Q_j(K), r_j(K)) - C_S^{\max}]$. It is easy to see that this latter expression is indeed the sum of the two preceding ones.

If backorders are allowed, the buyer can impose a stockout penalty p on the supplier in addition to consignment stock, in effect combining consignment (which revolves around *ownership* of inventory) with a vendor-managed inventory scheme (which revolves around *control* of inventory). The corresponding compensation payment becomes

$$\dot{P}_c^d(K) = -\frac{1}{2}\frac{hp}{h+p}\dot{Q}_c^d(K) = -\frac{1}{2}\frac{hp}{h+p}Q_c^d(K) + P_c^0.$$

The supplier will again choose $Q_c^d(K) = Q_j^d(K)$. All variable costs in the system and all decision rights have been allocated to the supplier in exchange for a fixed side payment. Essentially, the buyer has sold his firm to the supplier. Without knowing K, the buyer cannot determine the price the supplier would be willing to pay, which is why this approach is not optimal for the buyer. For the supplier, however, this raises the spectre of information asymmetry about p, a parameter that he will typically not know; we look at this in §7.5.

Consignment is often implemented for different commercial reasons. It could, for example, be part of a VMI scheme, to offer superior customer service, or the result of competitive pressure. We see here that consignment stock changes the incentive structure faced by the supplier and may as a result, even if unintended, improve supply-chain performance, precisely as happened in our case example.

To summarize, in this section we examined the situation without consignment, with the supplier setting Q, the buyer setting r, and the supplier holding private information about his setup cost K. Under full information, both parties can be better off by coordinating on the jointly optimal inventory policy. The buyer's optimal strategy is to coordinate on the joint optimum with a contract that is (just) acceptable to the supplier. Under asymmetric information, the buyer's optimal strategy is to coordinate on a jointly suboptimal inventory policy with a contract that is more attractive to the supplier as K decreases. The buyer could coordinate on the jointly optimal inventory policy implementing consignment stock, holding the supplier responsible for stockout costs, and offering a side payment; however, it is not optimal for the buyer to do so.

7. CONTRACTING ON *r* UNDER INFORMATION ASYMMETRY ABOUT *p*

In some ways, the situation in this section is the "reverse" of that considered above: Both parties have agreed to implement consignment stock, but the supplier wants the buyer to accept a lower level of consigned inventory. The supplier is now the party taking the initiative and therefore the principal, but the buyer holds private information about backorder cost p. We start from consignment game Γ^c, so the supplier incurs setup costs and holding costs; the buyer continues to incur backorder costs.

7.1. Coordinating on Reorder Point *r* Under Full Information

Here, too, coordinating on the jointly optimal reorder point r_j rather than letting the buyer insist on too high a reorder point, and sharing the benefits, can make both parties better off; i.e., $C_s(Q, r_j(Q)) - C_s(Q, r_b) \geqslant C_b(Q, r_j(Q)) - C_b(Q, r_b)\ \forall r_b$. The supplier could offer the buyer a contract $(Q, r_j(Q), P)$ specifying a compensation rate $P \in [C_b(Q, r_j(Q)) - C_b(Q, r_b), C_s(Q, r_b), -C_s(Q, r_j(Q))]$ in exchange for the buyer's accepting a reorder point $r_j(Q)$. P now depends on $C_b(\cdot)$ and therefore on the buyer's backorder cost p. Below we show how the supplier can offer a menu of contracts without knowing p.

7.2. The Supplier's Optimization Problem Under Asymmetric Information

Now only the buyer can observe p, and we follow a contracting approach analogous to that in the previous section; the supplier offers a menu of contracts $\{Q(\cdot), r(\cdot), P(\cdot)\}$, letting the buyer choose according to the latter's backorder cost p. The contracting process is analogous to that in the previous section. We also assume that there is no renegotiation after contracting, and that the supplier can verify and enforce the buyer's choice of $Q(\hat{p})$ and $r(\hat{p})$. This certainly holds for the production lot size because this is controlled by the supplier. It is less obvious for the reorder quantity, though, because that assumes the supplier has full information with respect to the buyer's inventory position. In the context of consignment stock, this is not unreasonable; however, the potential difficulties in enforcing the buyer's actual reordering behaviour immediately point to one of the pitfalls of consignment stock, which was also apparent in the case in §3. The supplier holds a prior distribution $F_p(\cdot)$, differentiable on its domain $[\underline{p}, \bar{p}]$ with $0 < \underline{p} < \bar{p} < \infty$ with density $f_p(\cdot)$, and such that $F_p(\cdot)$ satisfies the decreasing reversed hazard rate property (Appendix B), so that $D_p \frac{f_p(p)}{F_p(p)} \leqslant 0$. As before, we can derive the buyer's incentive-compatibility constraint and participation constraint (see Appendix C.1):

Incentive-compatibility constraint on $P(\cdot)$:

$$\textbf{IC:}\quad \dot{P}(p) = \frac{1}{Q}[G_p^-(r(p)+Q(p))(\dot{r}(p)+\dot{Q}(p)) - G_p^-(r(p))\dot{r}(p)] - \frac{\dot{Q}(p)}{Q(p)^2} \times \int_{r(p)}^{r(p)+Q(p)} G_p^-(y)\,dy \quad \forall p \in [\underline{p}, \bar{p}]. \qquad (24)$$

Participation constraint:

$$\textbf{PC:}\quad C_b^{\max} \geqslant C_b^c(Q(p), r(p)) - P(p) \quad \forall p \in [\underline{p}, \bar{p}]. \tag{25}$$

Again, we are not interested in the actual value of $C_b^{\max}$ because this will only shift $P(\cdot)$ up or down, but we focus on the shape of $P(\cdot)$. The supplier solves the following problem:

$$\mathscr{S}_2 \quad \min_{Q(\cdot), r(\cdot), P(\cdot)} E_{F_p}[P(p) + C_s^c(Q(p), r(p))], \qquad \text{s.t.} \quad \text{IC (24), PC (25).} \tag{26}$$

7.3. The Supplier's Optimal Menu of Contracts

Solving the optimal control problem $\mathscr{S}_2$ yields the optimal menu $\{Q(\cdot), r(\cdot), P(\cdot)\}$ shown below. We occasionally drop the dependence of $r(\cdot)$ and $Q(\cdot)$ on p when there is no danger of confusion.

Proposition 2. *Assume that the first-order conditions for $\mathscr{S}_2$ are also sufficient. Then the optimal menu of contracts $\{Q_2(\cdot), r_2(\cdot), P_2(\cdot)\}$, when the buyer holds private information about p, is such that*:

$$Q_2(p) = Q_j\left(p + \frac{F_p(p)}{f_p(p)}\right) \qquad \forall p \in [\underline{p}, \bar{p}], \tag{27}$$

$$G_h^+(r_2) + G^-_{p+\frac{F_p(p)}{f_p(p)}}(r_2) = G_h^+(r_2+Q_2) + G^-_{p+\frac{F_p(p)}{f_p(p)}}(r_2+Q_2) \qquad \forall p \in [\underline{p}, \bar{p}], \tag{28}$$

$$\dot{P}_2(p) = \frac{1}{Q_2(p)}[G_p^-(r_2+Q_2)(\dot{r}_2+\dot{Q}_2) - G_p^-(r_2)\dot{r}_2] - \frac{\dot{Q}_2}{Q_2^2}\int_{r_2}^{r_2+Q_2} G_p^-(y)\,dy \qquad \forall p \in [\underline{p}, \bar{p}]. \tag{29}$$

$Q_2(p) \geqslant Q_j(p)$, $r_2(p) > r_j(p)$, and $r_2(\cdot)$ is increasing in p, $P_2(\cdot)$ is decreasing in p. Moreover, the buyer's net costs after compensation, $C_b(Q_2(p), r_2(p)) - P_2(p)$, are decreasing in p.

The insights from Proposition 2 are, of course, analogous to those of Proposition 1. $(Q_2(\cdot), r_2(\cdot))$ is always less efficient than $(Q_j(\cdot), r_j(\cdot))$: $C_j(Q_2(p), r_2(p)) \geqslant C_j(Q_j(p), r_j(p))$ for all $p \in [\underline{p}, \bar{p}]$. At $p = \underline{p}$, i.e., the buyer has the lowest possible stockout cost p, $F_p(\underline{p}) = 0$, so that Equations (27) and (28) reduce to $Q_2(p) = Q_j(p)$ and $r_2(p) = r_j(p)$, respectively.

7.4. Base-Stock Model: Contracting to Reduce Safety Stock

If setup costs $K = 0$, we can simplify these expressions and relate them to the well-known newsboy formula. Clearly, for $K = 0$, the lot size $Q = 0$, too; and all inventory in the system is now caused by the demand uncertainty. One part of that inventory, of magnitude μ, is determined by the supplier's delivery lead time L; the remainder, $r_2 - \mu$, was defined as the safety stock. In that case, contracting on reorder point r reduces to the following.

Corollary 2. *Assume that $K = 0$ (so that $Q = 0$) and that the first-order conditions for $\mathscr{S}_2$ are sufficient. Then the optimal menu of contracts $\{r_2(\cdot), P_2(\cdot)\}$, when the buyer holds private information about p, is such that*

$$r_2(p) = \Psi^{-1}\left(\frac{p + \frac{F_p(p)}{f_p(p)}}{p + h + \frac{F_p(p)}{f_p(p)}}\right) \quad \forall p \in [\underline{p}, \bar{p}], \tag{30}$$

$$\frac{\dot{P}_2(p)}{\dot{r}_2(p)} = -p[1 - \Psi(r_2(p))] \quad \forall p \in [\underline{p}, \bar{p}]. \tag{31}$$

This helps us interpret the compensation function: The marginal decrease in the compensation rate per unit increase in the reorder point is the expected backorder cost savings $p[1 - \Psi(r_2(p))]$. The expected marginal savings to the buyer of announcing a higher p are exactly equal to the decrease in compensation he receives, removing his initiative to exaggerate p in order to obtain a higher consigned stock level than is jointly optimal. Coordination on this joint optimum could be achieved here by allocating all variable costs and all decision rights to the buyer, the party holding private information; in other words, if the supplier sells his firm to the buyer. Consignment stock is in fact moving in the exact opposite direction by allocating more costs (the holding costs) to the supplier, which explains why it exacerbates the impact of the information asymmetry here rather than reduce it as was the case in §6.

7.5. Contracting on Stockout Penalty

So far, we assumed the supplier did not pay the buyer a penalty in case of stockouts, and that the contract was limited to specifying a compensation payment and lot size and reorder point. In a true VMI setting, the parties would specify a stockout penalty V per unit per unit time in the contract, leaving the supplier to choose an inventory policy (Q, r). A priori, one might expect this to lead to a more efficient outcome because the supplier now internalizes stockout costs himself. However, we will see that this is not the case, and that this contract is in fact equivalent to the one in the previous subsection. Supplier and buyer cost functions become

$$C_s^c(Q(V), r(V), V) = \frac{\lambda K + \int_r^{r+Q}[G_h^+(y) + G_V^-(y)]\,dy}{Q}, \tag{32}$$

$$C_b^c(Q(V), r(V), V) = \frac{\int_r^{r+Q} G_{p-V}^-(y)\,dy}{Q}, \tag{33}$$

where $G_V^-(y)$ and $G_{p-V}^-(y)$ are expected backorder costs when unit backorder cost is V and $p - V$, respectively. Of course, $C_s^c(Q(V), r(V), V) + C_b^c(Q(V), r(V), V) = C_j(Q, r)$ for any V. (Although $r(V)$ is now actually chosen by the supplier after the buyer has chosen a $V(\hat{p})$, the problems of verifiability and enforceability remain, as the supplier may not be able to observe the buyer's inventory position and therefore be unable to verify the true nature of a declared stockout. This is precisely the problem observed in the case.) When the contracting process has determined V, the supplier will choose $r = r_j(V)$ and $Q = Q_j(V)$, i.e., an inventory policy that would be jointly optimal if true backorder costs were V. To find the optimal menu $\{V(\cdot), P(\cdot)\}$, the supplier solves optimal control problem $\mathcal{S}_3$ given in Appendix C.3.

PROPOSITION 3. *Assume the first-order conditions for $\mathcal{S}_3$ are also sufficient. Then the optimal menu of contracts $\{V_3(\cdot), P_3(\cdot)\}$, and the supplier's optimal inventory policy $(Q_3(V_3(\cdot)), r_3(V_3(\cdot)))$ when the buyer holds private information about p, is*

$$V_3(p) = p + \frac{F_p(p)}{f_p(p)} \quad \forall\, p \in [\underline{p}, \bar{p}], \tag{34}$$

$$\begin{aligned} \dot{P}_3(p) = &-\frac{\dot{Q}_3 \dot{V}_3}{Q_3^2} \int_{r_3}^{r_3+Q_3} G^-_{p-V_3}(y)\,dy \\ &+ \frac{\dot{V}_3}{Q_3}[G^-_{p-V}(r_3+Q_3)(\dot{r}_3+\dot{Q}_3) - G^-_{p-V}(r_3)\dot{r}_3] \\ &- \frac{\dot{V}_3}{Q_3} \int_{r_3}^{r_3+Q_3} G^-_1(y)\,dy \quad \forall\, p \in [\underline{p}, \bar{p}], \end{aligned} \tag{35}$$

$$Q_3(V_3(p)) = Q_2(p) = Q_j(V_3(p)) \quad \forall\, p \in [\underline{p}, \bar{p}], \tag{36}$$

$$r_3(p) = r_2(p) = r_j(V_3(p)) \quad \forall\, p \in [\underline{p}, \bar{p}]. \tag{37}$$

Buyer's and supplier's expected net costs are equal to those in Proposition 2. Moreover, $V_3(p) \geqslant p$ for all $p \in [\underline{p}, \bar{p}]$, and the overcompensation $(V_3(p) - p)$ is increasing in p.

This leads to two interesting observations. First, contracting on payment and stockout penalty V leads to the same stock level and expected costs for both parties as contracting on payment and inventory policy (Q, r). Second, the stockout penalty $V_3(p)$ paid by the supplier is always at least as great as the buyer's true backorder costs p. Under this contract, the buyer makes a *profit* on each stockout occasion. Although both menus $\{Q_2(\cdot), r_2(\cdot), P_2(\cdot)\}$ and $\{V_3(\cdot), P_3(\cdot)\}$ are equivalent (and the buyer stays below his reservation cost level) in expectation, the corresponding incentive structures are radically different: In the first, the buyer will attempt to avoid stockouts, whereas in the second, he has an incentive to *maximize* the number of stockouts. Explicit inclusion of such behavioral issues (which would add a "moral hazard" component to the current "adverse selection" model) is left for further research.

To summarize, in this section we examined the situation with consignment stock, with the supplier setting Q, the buyer setting r, and the buyer holding private information about stockout cost p. Under full information, both parties can coordinate on the jointly optimal inventory policy and can be better off than without coordination. The supplier's optimal strategy is to coordinate on the joint optimum with a contract that is (just) acceptable to the buyer. Under asymmetric information, the supplier's optimal strategy is to coordinate on a jointly suboptimal inventory policy with a contract that is more attractive to the buyer as p decreases. The supplier could coordinate on the jointly optimal inventory policy by passing holding costs back to the buyer (i.e., by ending the consignment stock arrangement), charging the buyer for the setup costs, and offering a side payment; however, it is not optimal for the supplier to do so. Supplier and buyer are indifferent between contracting directly on an inventory policy (Q, r) or indirectly through a stockout penalty V to be paid by the supplier to the buyer, letting the supplier determine Q and r himself based on the agreed value of V.

8. CONCLUSIONS AND FURTHER RESEARCH

This paper makes several theoretical contributions. It generalizes the framework for the (Q, r) model developed by Zheng (1992) to the case of a two-party supply chain with conflicting objectives and asymmetric information, and it examines how reallocating decision rights affects supply-chain performance. It also demonstrates how the *combined* presence of incentive conflicts and information asymmetries, typical of almost any supply chain, hurts overall performance. Specifically, it shows that consignment stock has an important incentive effect, sometimes mitigating and sometimes aggravating the resulting inefficiency.

We should repeat that achieving joint optimality is not difficult: In either case, a principal aiming for jointly optimal outcomes could make the other party bear all costs, in exchange for a fixed payment. The principal is selling the firm to the agent, but this is *not* optimal for the principal as the value of the firm to the agent depends on the agent's cost structure. Because this is unknown to the principal, he must sell the firm at a price low enough to satisfy all possible types of agents, which is highly unattractive. Alternatively, one could assume the principal also has an outside reservation cost level and give him the option of not dealing with agents that would leave him in excess of that reservation cost level. The shape of the resulting optimal contracts for agents that are not excluded is as before; the entire compensation payment shifts downward to make the participation constraint binding for the "worst" nonexcluded agent. In either case, the principal will *not* choose any mechanism that optimizes the entire supply chain. This is what sets the current framework apart from most previous work. We have shown which mechanism the principal could use to enforce joint optimality, but we also derived what the principal's *optimal* mechanism would be. In the absence of a central planner with full information, no party can induce jointly optimal behavior from all parties in the supply chain without sacrificing his own profits.

The findings reported here clearly show that ignoring incentive and information issues can lead to undesirable behavior. The practical implications are significant. In (re-)designing a supply chain, it is tempting to take the perspective of a central planner and focus on improving overall efficiency. Few supply chains have anything that can pass for a central planner, so this is rarely an option. Instead, if possible, each party should be made to fully internalize all consequences of his decisions, and should have all information available to support those decisions. (This principle holds only when all parties are risk-neutral, but in a supply-chain inventory management context this seems reasonable.) Specifically, this paper proposes a framework for how incentives should be structured to help reduce supply-chain inventory. If long supplier production cycles are the key driver of inventory, the supplier should be made to bear the costs of the resulting cycle stocks, e.g., through a consignment scheme. Conversely, if uncertainty about downstream demand is the main driver, the buyer should bear the costs of the resulting safety stock. This principle can, of course, be extended to any other cause of inventory, such as yield uncertainty, transit times, etc.

This work immediately points to several avenues for future research. Including more levels of the supply chain and also considering, e.g., in-transit and other inventories could help understand supply-chain performance. Competition may be another way to overcome information asymmetries, but poses other dilemmas. The current paper focused on a very basic incentive conflict in inventory management under asymmetric information; ongoing work (Corbett and DeCroix 2001) studies a similarly fundamental incentive conflict related to amount of effort supplier and buyer exert to improve supply chain efficiency.

ACKNOWLEDGMENTS

This paper is partly based on the author's doctoral dissertation submitted at INSEAD and completed while visiting Vanderbilt University; financial assistance from these institutions is gratefully acknowledged. The author thanks seminar participants at UCLA, Stanford, Northwestern, Wharton, Columbia, and several conferences and workshops for their helpful comments. Two anonymous referees and the associate editor provided very valuable and constructive feedback that led to substantial improvements in the paper. The author is especially grateful to Teck Ho, Marty Larivière, Christoph Loch, and Luk Van Wassenhove for their suggestions, and above all to Xavier de Groote, who sadly died before this work was completed, for having provided much of the inspiration for it.

APPENDIX A: PROOFS FROM SECTION 4: LEMMA 1

The statements correspond to Lemma 2 in Zheng (1992, p. 90), Lemma 3.1 in Zheng, Lemma 3.3 in Zheng, Proposition 1 in Zheng and de Groote (1993, p. 4), Lemma 6 in Zheng, Lemma 3.4 in Zheng, and expressions (20) and (21) in Zheng (1992, p. 94), respectively.

APPENDIX B: THE REVERSED HAZARD RATE

The assumption that $D_x \frac{f(x)}{F(x)} \leqslant 0$ for x defined on $[\underline{x}, \bar{x}]$ with $0 \leqslant \underline{x} < \bar{x} < \infty$ is often referred to in the economics literature as the monotone decreasing hazard rate property (and expressed as $D_x \frac{F(x)}{f(x)} \geqslant 0$); see, for example, Laffont and Tirole (1993, p. 66). Strictly speaking, though, the hazard rate of a distribution $F(x)$ is defined as $\frac{f(x)}{1-F(x)}$ (see, e.g., Barlow and Proschan 1975). The reversed hazard rate is defined as $\frac{f(x)}{F(x)}$ in Shaked and Shanthikumar (1994, p. 24). The decreasing reversed hazard rate assumption is also equivalent to requiring log-concavity of F, i.e., that log F be concave on its domain (Bagnoli and Bergstrom 1989). Many common distributions satisfy this decreasing reversed hazard rate assumption, including uniform, normal, logistic, chi-squared, and exponential. We do have to assume that the support of $f(x)$ is finite because the standard solution procedure for the associated contracting optimization problem (see Appendix C) relies on setting the participation constraint to be binding at $x = \bar{x}$. However, it is easy to verify Theorem 7 in Bagnoli and Bergstrom (1989, p. 10), which assures us that any finite truncation of any infinite distribution satisfying this decreasing reversed hazard rate property also satisfies the same property.

APPENDIX C: PROOFS FROM SECTION 7

C.1. Incentive-Compatibility Constraint

To derive (24), we follow Laffont and Tirole (1993) and first formulate the buyer's optimization problem of deciding which $\hat{p}$ to announce when offered a menu $\{Q(\cdot), r(\cdot), P(\cdot)\}$:

$$\mathscr{B}\ \max_{\hat{p}} \Pi_b(p, \hat{p}) := P(\hat{p}) - \frac{1}{Q(\hat{p})} \int_{r(\hat{p})}^{r(\hat{p})+Q(\hat{p})} G_p^-(y)\,dy. \quad (38)$$

The first-order condition for $\mathscr{B}$ is

$$\dot{P}(\hat{p}) = -\frac{\dot{Q}(\hat{p})}{Q(\hat{p})^2} \int_{r(\hat{p})}^{r(\hat{p})+Q(\hat{p})} G_p^-(y)\,dy + \frac{1}{Q(\hat{p})}[G_p^-(r(\hat{p}) + Q(\hat{p}))(\dot{r}(\hat{p}) + \dot{Q}(\hat{p})) - G_p^-(r(\hat{p}))\dot{r}(\hat{p})]. \quad (39)$$

Requiring that this be solved at $\hat{p} = p$ yields the desired result.

The proof that this first-order condition (39) is sufficient for the buyer's optimization problem is analogous to that in Laffont and Tirole (1993, p. 121). The proof will follow by contradiction. Assume there is a $\hat{p} \neq p$ such that $\Pi_b(p, \hat{p}) > \Pi_b(p, p)$. This is equivalent to $\int_p^{\hat{p}} D_2\Pi_b(p, x)dx > 0$. From first-order condition (39), we know that $D_2\Pi_b(x, x) = 0$ for all x, so that the previous expression is equivalent to $\int_p^{\hat{p}}[D_2\Pi_b(p, x) - D_2\Pi_b(x, x)]dx > 0$, which in turn is equivalent to $\int_p^{\hat{p}} \int_x^p D_{21}\Pi_b(u, x)\,du\,dx > 0$. One can easily verify that $D_{21}\Pi_b(u, x) = \frac{1}{uQ(x)}[\frac{\dot{Q}(x)}{Q(x)^2} \int_{r(x)}^{r(x)+Q(x)} G_u^-(y)dy - (G_u^-(r(x) + Q(x))(\dot{r}(x) + \dot{Q}(x)) - G_u^-(r)\dot{r}(x))]$. Because $\dot{r}(p) \geqslant 0$ in the optimal menu of contracts (see below) and $G_u^-(y)$ is decreasing in y for any u, we know that $\dot{r}(x)G_u^-(r(x) + Q(x)) \leqslant \dot{r}(x)G_u^-(r(x))$. Therefore, $D_{21}\Pi_b(u, x) \geqslant 0$ if $\frac{\dot{Q}(x)}{Q(x)}[\frac{1}{Q(x)} \int_{r(x)}^{r(x)+Q(x)} G_u^-(y)dy - G_u^-(r(x) + Q(x))] \geqslant 0$. The second factor is nonnegative as $G_u^-(\cdot)$ is decreasing in y. So, if $\hat{p} > p$, then $x > p$ for all $x \in [p, \hat{p}]$, so that the double integral above is nonpositive, which is a contradiction. A similar contradiction follows for the case $\hat{p} < p$, which completes the proof. □

C.2. Proof of Proposition 2: Solving the Supplier's Optimization Problem

Following, for example, Kamien and Schwartz (1981), we construct the Hamiltonian for $\mathscr{S}_1$. Write $v(\cdot)$ for $\dot{Q}(\cdot)$, $u(\cdot)$ for $\dot{r}(\cdot)$, and formulate $\mathscr{S}_1$ as a maximization problem:

$$\begin{aligned} &H(Q, r, P, u, v, \lambda_Q, \lambda_r, \lambda_P) \\ &\quad = -\left[P(p) + \frac{\lambda K + \int_r^{r+Q} G_h^+(y)\,dy}{Q}\right] f(p) \\ &\qquad + \lambda_P \frac{1}{Q}[G_p^-(r+Q)(u+v) - G_p^-(r)u] \\ &\qquad - \lambda_P \frac{v}{Q^2} \int_r^{r+Q} G_p^-(y)\,dy + \lambda_Q v + \lambda_r u. \end{aligned} \tag{40}$$

An optimum must satisfy the following conditions:

$$\dot{\lambda}_Q = -D_Q H, \tag{41}$$

$$\dot{\lambda}_r = -D_r H, \tag{42}$$

$$\dot{\lambda}_P = -D_P H, \tag{43}$$

$$D_u H = 0, \tag{44}$$

$$D_u H = 0. \tag{45}$$

The boundary $p = \underline{p}$ is unconstrained, so the transversality condition on $\overline{\lambda}_P(\underline{p})$ becomes $\lambda_P(\underline{p}) = 0$. This and Equation (43), which reduces to $\dot{\lambda}_P(p) = f(p)$, give $\lambda_P(\cdot) = F_p(\cdot)$. From that and (44) we find $\lambda_r(p) = -\frac{F_p(p)}{Q}[G_p^-(r+Q) - G_p^-(r)]$. Differentiating this and equating it to (42), recalling that $u(\cdot) = \dot{r}(\cdot)$, yields the desired result for $r_2(\cdot)$. Similarly, from (45) we find $\lambda_Q = -\frac{F_p(p)}{Q} G_p^-(r+Q) + \frac{F_p(p)}{Q^2} \int_r^{r+Q} G_p^-(y)\,dy$. Differentiating and equating to (41) yields the expression for $Q_2(\cdot)$. The expression for $\dot{P}_2(\cdot)$ follows from (24).

So far we have ignored the participation constraint (25). We can set $P_2(\bar{p})$ such that (25) is binding for $p = \bar{p}$; we need to check whether this implies that (25) is satisfied for all $p \leqslant \bar{p}$. A sufficient condition for this is that $C_b(p, Q_2(p), r_2(p)) - P_2(p)$ be increasing in p. Because $D_p C_b(p, Q_2(p), r_2(p)) - \dot{P}_2(p) = \frac{1}{pQ_2(p)} \times \int_{r_2(p)}^{r_2(p)+Q_2(p)} G_p^-(y)\,dy \geqslant 0$, this condition is met.

The necessary first-order conditions yield a unique stationary point, which by the assumption in the proposition must be the true optimum. The inequalities in the proposition are easily verified. To see that $r_2(p)$ is increasing in p (which we used above to ensure sufficiency of the buyer's first-order condition), write r_2 as a function of p and Q_2. Then $r_2(p, Q_2(p)) = r_j(p + \frac{F_p(p)}{f_p(p)}, Q_j(p + \frac{F_p(p)}{f_p(p)}))$, so by the decreasing reversed hazard rate assumption and Zheng and de Groote's (1993) observation that $r_j(p, Q_j(p))$ is increasing in p, a fortiori $r_2(p) = r_j(p + \frac{F_p(p)}{f_p(p)}, Q_j(p + \frac{F_p(p)}{f_p(p)}))$ must also be increasing in p. (Laffont and Tirole 1993 make additional assumptions, including one that the third derivative of their "disutility of effort" function be nonnegative, which allow them to dispense with the assumption that the first-order conditions are also sufficient. This latter assumption is quite common in the economics literature, as sufficiency is frequently impossible to prove. Moreover, as our results are fully analogous to those from Laffont and Tirole 1993 and elsewhere, the sufficiency assumption seems reasonable here too.) □

C.3. Contracting on Stockout Penalty

Formulation and solution of problem $\mathscr{S}_3$, yielding the result in Proposition 3, are largely analogous to those of Proposition 2. The supplier solves $\mathscr{S}_3$

$$\mathscr{S}_3 \quad \min_{\substack{V(\cdot),\, P(\cdot) \\ \text{s.t.}}} E_F[P(p) + C_s^c(Q(V), r(V), V(p))], \tag{46}$$

$$\begin{aligned} \textbf{IC: } \dot{P}(p) = &-\frac{\dot{Q}\dot{V}}{Q^2} \int_r^{r+Q} G_{p-V}^-(y)\,dy \\ &+ \frac{\dot{V}}{Q}[G_{p-V}^-(r+Q)(\dot{r}+\dot{Q}) - G_{p-V}^-(r)\dot{r}] \\ &- \frac{\dot{V}}{Q} \int_r^{r+Q} G_1^-(y)\,dy \quad \forall p \in [\underline{p}, \bar{p}], \end{aligned} \tag{47}$$

$$\begin{aligned} \textbf{PC: } &C_b^c(Q(V), r(V), V(p)) \\ &\quad - P(p) \leqslant C_b^{\max} \quad \forall p \in [\underline{p}, \bar{p}]. \end{aligned} \tag{48}$$

After tedious computation, this yields the first-order condition that

$$A_1 - A_2 + A_3 = 0, \tag{49}$$

where

$$\begin{aligned} A_1 = &f_p(p) \frac{\lambda K}{Q^2} \dot{Q} + f_p(p) \frac{\dot{Q}}{Q^2} \int_r^{r+Q} [G_h^+(y) + G_V^-(y)]\,dy \\ &+ f_p(p) \frac{\dot{Q}}{Q} \big[G_h^+(r+Q) + G_V^-(r+Q)\big] \\ &\forall p \in [\underline{p}, \bar{p}], \end{aligned} \tag{50}$$

$$\begin{aligned} A_2 = &f_p(p) \frac{\dot{r}}{Q}[G_h^+(r+Q) + G_V^-(r+Q) \\ &- G_h^+(r) - G_V^-(r)] \quad \forall p \in [\underline{p}, \bar{p}], \end{aligned} \tag{51}$$

$$\begin{aligned} A_3 = &\frac{\dot{Q}}{Q^2} \int_r^{r+Q} G_{f_p(p)(p-V)+F_p(p)}(y)\,dy \\ &- \frac{1}{Q}[G_{f_p(p)(p-V)+F_p(p)}(r+Q)(\dot{r}+\dot{Q}) \\ &- G_{f_p(p)(p-V)+F_p(p)}(r)\dot{r}] \quad \forall p \in [\underline{p}, \bar{p}]. \end{aligned} \tag{52}$$

Because the supplier chooses (Q, r) to be the joint optimum with backorder costs V, we know that $A_1 \equiv 0$ and $A_2 \equiv 0$ for all V, as these are exactly the joint optimality conditions. Therefore, we must have $A_3 \equiv 0$, for any choice of p and K and with (Q, r) defined by A_1 and A_2. We know that $G^-_\beta(y) \geqslant 0\ \forall \beta$, so we must have $G_{f_p(p)(p-V)+F_p(p)}(y) \equiv 0$ for all y. Because $G_{f_p(p)(p-V)+F_p(p)}(y) = [f_p(p)(p-V)+F_p(p)]G^-_1(y)$ and $G^-_1(y) > 0$, this yields the desired result that $f_p(p)(p-V_3(p))+F_p(p) \equiv 0$.

To verify equivalence of $\mathscr{S}_2$ and $\mathscr{S}_3$, observe first that the buyer's net costs (after compensation) are equal in both cases: $C^c_b(Q_3(V_3(\bar{p})), r_3(V_3(\bar{p}))) - P_3(\bar{p}) = C^c_b(Q_2(\bar{p}), r_2(\bar{p})) - P_2(\bar{p})$ and $D_p[C^c_b(Q_3(V_3(p)), r_3(V_3(p))) - P_3(p)] = D_p[C^c_b(Q_2(p), r_2(p)) - P_2(p)]$. Here we use the fact that $D_pQ_3(V_3(p)) = \dot{Q}_3(V_3)\dot{V}_3(p)$ and similarly for $r_3(\cdot)$. Total costs are equal in both cases (because the inventory policies are identical), so the supplier's net costs must be equal in both cases too. □

REFERENCES

Axsäter, S., K. Rosling. 1993. Installation vs. echelon stock policies for multilevel inventory control. *Management Sci.* **39** 1274–1280.

Bagnoli, M., T. Bergstrom. 1989. Log-concave probability and its applications. Working Paper, University of Michigan, October.

Barlow, R. E., F. Proschan. 1975. *Statistical Theory of Reliability and Life Testing*. Holt, Rinchart & Winston, New York.

Browne, S., P. Zipkin. 1989. Inventory models with continuous, stochastic demands. Working paper, Graduate School of Business, Columbia University.

Cachon, G. P. 1997. Stock wars: inventory competition in a two echelon supply chain with multiple retailers. Working Paper, Fuqua School of Business, Duke University.

——, M. A. Larivière. 1999. Capacity choice and allocation: strategic behaviour and supply chain performance. *Management Sci.* **45** 1091–1108.

——, ——. 1999. Contracting to assure supply: How to share demand forecasts in a supply chain. Working Paper, Fuqua School of Business, Duke University.

——, P. Zipkin. 1999. Competitive and cooperative inventory policies in a two-stage supply chain. *Management Sci.* **45** 936–953.

Chen, F. 1999. Decentralized supply chains subject to information delays. *Management Sci.* **45** 1076–1090.

Clark, A. J., H. Scarf. 1960. Optimal policies for a multi-echelon inventory problem. *Management Sci.* **6** 475–490.

Corbett, C. J. 1996. Supply chain logistics in an imperfect partnership (Inventory management and incentive and information asymmetries). Unpublished Ph.D. dissertation, INSEAD, Fontainebleau, France.

——, J. D. Blackburn, L. N. Van Wassenhove. 1999. Partnerships to improve supply chains. *Sloan Management Review* **40**(Summer) 71–82.

——, G. A. DeCroix. 2001. Shared savings contracts for indirect materials in supply chains: Channel profits and environmental impacts. *Management Sci.* **47**(7).

——, X. de Groote. 1997. Integrated supply-chain lot sizing under asymmetric information. Manuscript, The Anderson School at UCLA; presented at First Xavier de Groote Memorial conference, INSEAD, Fontainebleau, France, May 16, 1997.

——, ——. 2000. A supplier's optimal quantity discount policy under asymmetric information. *Management Sci.*, **46** 444-450.

Dada, M., K. N. Srikanth. 1987. Pricing policies for quantity discounts. *Management Sci.* **33** 1247–1252.

Fudenberg, D., J. Tirole. 1991. *Game Theory*. The MIT Press, Cambridge, MA.

Ha, A. 1997. Supply contract for short-life-cycle product with demand uncertainty and asymmetric cost information. Working Paper, Yale School of Management.

Kamien, M. I., N. L. Schwartz. 1981. *Dynamic Optimization (The Calculus of Variations and Optimal Control in Economics and Management)*. North-Holland, New York.

Laffont, J.-J., J. Tirole. 1993. *A Theory of Incentives in Procurement and Regulation*. The MIT Press, Cambridge, MA.

Lal, R., R. Staelin. 1984. An approach for developing an optimal discount pricing policy. *Management Sci.* **30** 1524–1539.

Lee, H. L., M. J. Rosenblatt. 1986. A generalized quantity discount pricing model to increase supplier's profits. *Management Sci.* **32** 1177–1185.

——, S. Whang. 1996. Decentralized multi-echelon supply chains: incentives and information. Working Paper, Stanford University.

Monahan, J. P. 1984. A quantity discount pricing model to increase vendor profits. *Management Sci.* **30** 720–726.

Myerson, R. 1979. Incentive compatibility and the bargaining problem. *Econometrica* **47** 61–74.

Serfozo, R., S. Stidham. 1978. Semi-stationary clearing processes. *Stochastic Process. Appl.* **6** 165–178.

Shaked, M., J. G. Shanthikumar. 1994. *Stochastic Orders and Their Applications*. Academic Press, Inc., San Diego, CA.

Tsay, A. A., S. Nahmias, N. Agrawal. 1999. Modeling supply chain contracts. A review. Chapter 10 in *Quantitative Models for Supply Chain Management*. S. Tayur, M. Magazine, R. Ganeshan (eds.). Kluwer Academic Publishers.

Weng, Z. K. 1995. Channel coordination and duantity discounts. *Management Sci.* **41** 1509–1522.

Zheng, Y.-S. 1992. On properties of stochastic inventory systems. *Management Sci.* **38** 87–103.

——, X. de Groote. 1993. A sensitivity analysis of stochastic inventory systems. Revised version of INSEAD Working Paper 92/16/TM.