

Coordination of Supply Chains with Risk-Averse Agents

Xianghua Gan • Suresh P. Sethi • Houmin Yan

Department of Information Systems and Operations Management,
The University of Texas at Dallas,
Dallas, Texas 75083, USA

The extant supply chain management literature has not addressed the issue of coordination in supply chains involving risk-averse agents. We take up this issue and begin with defining a coordinating contract as one that results in a Pareto-optimal solution acceptable to each agent. Our definition generalizes the standard one in the risk-neutral case. We then develop coordinating contracts in three specific cases: (i) the supplier is risk neutral and the retailer maximizes his expected profit subject to a downside risk constraint; (ii) the supplier and the retailer each maximizes his own mean-variance trade-off; and (iii) the supplier and the retailer each maximizes his own expected utility. Moreover, in case (iii), we show that our contract yields the Nash Bargaining solution. In each case, we show how we can find the set of Pareto-optimal solutions, and then design a contract to achieve the solutions. We also exhibit a case in which we obtain Pareto-optimal sharing rules explicitly, and outline a procedure to obtain Pareto-optimal solutions.

Key words: supply chain management; Pareto-optimality; coordination; risk aversion; Nash Bargaining solution

Submissions and Acceptance: Received xxxx xxxx; revision received xxxx xxxx; accepted xxxx xxxx.

1. Introduction

Much of the research on decision making in a supply chain has assumed that the agents in the supply chain are risk neutral, i.e., they maximize their respective expected profits. An important focus of this research has been the design of supply contracts that coordinate the supply chain. When each of the agents maximizes his expected profit, the objective of the supply chain considered as a single entity is unambiguously to maximize its total expected profit. This fact alone makes it natural to define a supply chain to be coordinated if the chain's expected profit is maximized and each agent's reservation profit is met. A similar argument holds if each agent's objective is to minimize his expected cost.

In this paper, we consider supply chains with risk-averse agents. Simply put, an agent is risk averse if the agent prefers a certain profit π to a risky profit, whose expected value equals π . In the literature, there are many measures of risk aversion; see Szegö (2004) for examples. Regardless of the measure used, when one or more agents in the supply chain are risk averse, it is

no longer obvious as to what the objective function of the supply chain entity should be. Not surprisingly, the issue of coordination of supply chain consisting of risk-averse agents has not been studied in the supply chain management literature. That is not to say that the literature does not realize the importance of the risk-averse criteria. Indeed, there are a number of papers devoted to the study of inventory decisions of a single risk-averse agent. These include Lau (1980), Bouakiz and Sobel (1992), Eeckhoudt, Gollier, and Schlesinger (1995), Chen and Federgruen (2000), Agrawal and Seshadri (2000a), Buzacott, Yan, and Zhang (2001), Chen, Sim, Simchi-Levi, and Sun (2003), and Gaur and Seshadri (2003). There also have been a few studies of supply chains consisting of one or more risk-averse agents. Lau and Lau (1999) and Tsay (2002) consider decision making by a risk-averse supplier and a risk-averse retailer constituting a supply chain. Agrawal and Seshadri (2000b) introduce a risk-neutral intermediary to make ordering decisions for risk-averse retailers, whose respective profits are side payments from the intermediary. Van Mieghem (2003)

has reviewed the literature that incorporates risk aversion in capacity investment decisions. While these papers consider risk-averse decision makers by themselves or as agents in a supply chain, they do not deal with the issue of the supply chain coordination involving risk-averse agents.

It is this issue of coordination of supply chains consisting of one or more risk-averse agents that is the focus of this paper. That many decision makers are risk-averse has been amply documented in the finance and economics literature; see, for example, von Neumann and Morgenstern (1944), Markowitz (1959), Jorion (2001), and Szegö (2004). We shall therefore develop the concept of what we mean by coordination of a supply chain, and then design explicit contracts that achieve the defined coordination.

For this purpose, we use the Pareto-optimality criterion, used widely in the group decision theory, to evaluate a supply chain's performance. We define each agent's payoff to be a real-valued function of a random variable representing his profit, and propose that a supply chain can be treated as coordinated if no agent's payoff can be improved without impairing someone else's payoff and each agent receives at least his reservation payoff. We consider three specific cases of a supply chain: (i) the supplier is risk neutral and the retailer maximizes his expected profit subject to a downside risk constraint, (ii) the supplier and the retailer each maximizes his own mean-variance trade-off, and (iii) the supplier and the retailer each maximizes his own expected utility. We show how we can coordinate the supply chain in each case according to our definition. In each case, we do this by finding the set of Pareto-optimal solutions acceptable to each agent, and then constructing a flexible contract that can attain any of these solutions. Moreover, the concept we develop and the contracts we obtain generalize the same known for supply chains involving risk-neutral agents.

The remainder of the paper is organized as follows. In Section 2, we review the related literature in supply chain management and group decision theory. In Section 3, we introduce a definition of coordination of a supply chain consisting of risk-averse agents. In Section 4, we characterize the Pareto-optimal solutions and find coordinating contracts for the supply chains listed as the first two cases. In Section 5, we first take up the third case using exponential utility functions for the agents, and design coordinating contracts as well as obtain the Nash Bargaining solution. Then, we examine a case in which the supplier has an exponential utility followed by a linear utility. Section 6 provides a discussion of our results. The paper concludes in Section 7 with suggestions for future research.

2. Literature Review

There is a considerable literature devoted to contracts that coordinate a supply chain involving risk-neutral agents. This literature has been surveyed by Cachon (2002). In addition, the book by Tayur, Ganeshan, and Magazine (1999) contains a number of chapters addressing supply contracts. In light of these, we limit ourselves to reviewing papers studying inventory and supply chain decisions by risk-averse agents. First, we review papers dealing with a single risk-averse agent's optimal inventory decision. Then, we review articles dealing with decision making by risk-averse agents in a supply chain.

Chen and Federgruen (2000) re-visit a number of basic inventory models using a mean-variance approach. They exhibit how a systematic mean-variance trade-off analysis can be carried out efficiently, and how the resulting strategies differ from those obtained in the standard analyses.

Agrawal and Seshadri (2000a) consider how a risk-averse retailer, whose utility function is increasing and concave in wealth, chooses the order quantity and the selling price in a single-period inventory model. They consider two different ways in which the price affects the distribution of demand. In the first model, they assume that a change in the price affects the scale of the distribution. In the second model, a change in the price only affects the location of the distribution. They show that, in comparison to a risk-neutral retailer, a risk-averse retailer will charge a higher price and order less in the first model, whereas he will charge a lower price in the second model.

Buzacott, Yan, and Zhang (2001) model a commitment and option contract for a risk-averse newsvendor with a mean-variance objective. The contract, also known as a take-or-pay contract, belongs to a class of volume flexible contracts, where the newsvendor reserves a capacity with initial information and adjusts the purchase at a later stage when some new information becomes available. They compare the performance of strategies developed for risk-averse and risk-neutral objectives. They conclude that the risk-averse objective can be an effective approach when the quality of information revision is not high. Their study indicates that it is possible to reduce the risk (measured by the variance of the profit) by six- to eight-fold, while the loss in the expected profit is almost invisible. On the other hand, the strategy developed for the expected profit objective can only be considered when the quality of information revision is high. They show furthermore that these findings continue to hold in the expected utility framework. The paper points out a need for modeling approaches that deal with downside risk considerations.

Lau and Lau (1999) study a supply chain consisting

of a monopolistic supplier and a retailer. The supplier and the retailer employ a return policy, and each of them has a mean-variance objective function. Lau and Lau obtain the optimal wholesale price and return credit for the supplier to maximize his utility. However, they do not consider the issue of improving the supply chain's performance, i.e., improving both players' utilities.

Agrawal and Seshadri (2000b) consider a single-period model in which multiple risk-averse retailers purchase a single product from a common supplier. They introduce a risk-neutral intermediary into the channel, who purchases goods from the vendor and sells them to the retailers. They demonstrate that the intermediary, referred to as the distributor, orders the optimal newsvendor quantity from the supplier and offers a menu of mutually beneficial contracts to the retailers. In every contract in the menu, the retailer receives a fixed side payment, while the distributor is responsible for the ordering decisions of the retailers and receives all their revenues. The menu of contracts simultaneously: (i) induces every risk-averse agent to select a unique contract from it; (ii) maximizes the distributor's profit; and (iii) raises the order quantities of the retailers to the expected value maximizing (newsvendor) quantities.

Tsay (2002) studies how risk aversion affects both sides of the supplier–retailer relationship under various scenarios of relative strategic power, and how these dynamics are altered by the introduction of a return policy. The sequence of play is as follows: first the supplier announces a return policy, and then the retailer chooses order quantity without knowing the demand. After observing the demand, the retailer chooses the price and executes on any relevant terms of the distribution policy as appropriate (e.g., returning any overstock as allowed). Tsay shows that the behavior under risk aversion is qualitatively different from that under risk neutrality. He also shows that the penalty for errors in estimating a channel partner's risk aversion can be substantial.

In a companion paper (Gan, Sethi, and Yan 2003), we examine coordinating contracts for a supply chain consisting of one risk-neutral supplier and one downside-risk-averse retailer. There, we design an easy-to-implement risk-sharing contract that accomplishes the coordination as defined in this paper.

Among these supply chain papers, Lau and Lau (1999) and Tsay (2002) consider the situation in which both the retailer and the supplier in the channel are risk averse. However, neither considers the issue of the Pareto-optimality of the actions of the agents. The aim of Agrawal and Seshadri (2000b) is to design a contract that increases the channel's order quantity to the optimal level in the risk-neutral case by having the risk-neutral agent assume all the risk. Once again, they

do not mention the Pareto-optimality aspect of the decision they obtain.

Finally, since our definition of coordination is based on the concepts used in the group decision theory, we briefly review this stream of literature. From the early fifties to the early eighties, a number of papers and books appeared that deal with situations in which a group faces intertwined external and internal problems. The external problem involves the choice of an action to be taken by the group, and the internal problem involves the distribution of the group payoff among the members. Arrow (1951) conducted one of the earliest studies on the group decision theory, and showed that, given an ordering of consequences by a number of individuals, no group ordering of these consequences exists that satisfies a set of seemingly reasonable behavioral assumptions. Harsanyi (1955) presented conditions under which the total group utility can be expressed as a linear combination of individuals' cardinal utilities. Wilson (1968) used Pareto-optimality as the decision criterion and constructed a group utility function to find Pareto-optimal solutions. Raiffa (1970) illustrates the criterion of Pareto-optimality quite lucidly, and discusses how to choose a Pareto-optimal solution in bargaining and arbitration problems. LaValle (1978) uses an allocation function to define Pareto-optimality. Eliashberg and Winkler (1981) investigate properties of sharing rules and the group utility functions in additive and multilinear cases.

3. Definition of Coordination of a Supply Chain with Risk-Neutral or Risk-Averse Agents

In this section, we define coordination of a supply chain consisting of agents that are risk neutral or risk averse. We use concepts developed in group decision theory that deals with situations in which a group faces intertwined external and internal problems. The external problem involves the choice of an action to be taken by the group, and the internal problem involves the distribution of the group payoff among the members. In group decision problems, a joint action of the group members is said to be *Pareto-optimal* if there does not exist an alternative action that is at least as acceptable to all and definitely preferred by some. In other words, a joint action is Pareto-optimal if it is not possible to make one agent better off without making another one worse off. We call the collection of all Pareto-optimal actions the *Pareto-optimal set*. It would not be reasonable for the group of agents to choose a joint action that is not Pareto-optimal. Raiffa (1970) and LaValle (1978) illustrate this idea quite lucidly with a series of examples.

A supply chain problem is obviously a group deci-

sion problem. The channel faces an external problem and an internal problem. External problems include decisions regarding order/production quantities, item prices, etc. The internal problem is to allocate profit by setting the wholesale price, deciding the amount of a side payment, if any, refund on the returned units, etc. Naturally, we can adopt the Pareto-optimality criterion of the group decision theory for making decisions in a supply chain. Indeed, in the risk-neutral case, the optimal action under a coordinating contract is clearly Pareto-optimal. In general, since the agents in the channel would not choose an action that is not in the Pareto-optimal set, the first step to coordinate a channel is to characterize the set. Following the ideas of Raiffa (1970) and LaValle (1978), we formalize below the definition of Pareto-optimality.

Let (Ω, \mathcal{F}, P) denote the probability space and N denote the number of agents in the supply chain, $N \geq 2$. Let S_i be the external action space of agent i , $i = 1, \dots, N$, and $S = S_1 \times \dots \times S_N$. For any given external joint action $s = (s_1, \dots, s_N) \in S$, the channel's total profit is a random variable $\Pi(s, \omega)$, $\omega \in \Omega$. Let E and V denote the expectation and variance operators defined on (Ω, \mathcal{F}, P) , respectively.

Now we define a sharing rule that governs the splitting of the channel profit among the agents. Let Θ be the set of all functions from $S \times \Omega$ to R^N .

DEFINITION 3.1. A function $\theta(s, \omega) \in \Theta$ is called a sharing rule if $\sum_i \theta_i(s, \omega) = 1$ almost surely. Under the sharing rule $\theta(s, \omega)$, agent i 's profit is represented by

$$\Pi_i(s, \omega, \theta(s, \omega)) = \theta_i(s, \omega)\Pi(s, \omega), i = 1, \dots, N.$$

Often, when there is no confusion, we write $\Pi(s, \omega)$ simply as $\Pi(s)$, $\theta(s, \omega)$ as $\theta(s)$, and $\Pi_i(s, \omega, \theta(s, \omega))$ as $\Pi_i(s, \theta(s))$. A supply chain's external problem is to choose an $s \in S$ and its internal problem is to choose a function $\theta(s) \in \Theta$. Thus the channel's total problem is to choose a pair $(s, \theta(s)) \in S \times \Theta$.

Now we define the preferences of the agents over their random profits. Let Γ denote the space of all random variables defined on (Ω, \mathcal{F}, P) . For $X, X' \in \Gamma$, the agent i 's preference will be denoted by a real-valued payoff function $u_i(\cdot)$ defined on Γ . The relation $u_i(X) > u_i(X')$, $u_i(X) < u_i(X')$, and $u_i(X) = u_i(X')$ indicate X is preferred to, less preferred to, and equivalent to X' , respectively.

It should be noted that this definition of payoff function allows for ordinal as well as cardinal utility functions. We provide following examples of payoff functions.

EXAMPLE 3.1. If agent i wants to maximize his mean-variance trade-off, then his payoff function is $u_i(X) = E(X) - \lambda V(X)$, $X \in \Gamma$, for some $\lambda > 0$.

EXAMPLE 3.2. Assume that agent i maximizes his expected profit under the constraint that the probabil-

ity of his profit being less than his target profit level α does not exceed a given level β , $0 < \beta \leq 1$. Then his payoff u_i can be represented as

$$u_i(X) = \begin{cases} E(X), & \text{if } P(X \leq \alpha) \leq \beta, \\ -\infty, & \text{if } P(X \leq \alpha) > \beta. \end{cases}$$

EXAMPLE 3.3. Suppose agent i has a concave increasing utility function $g_i: R^1 \rightarrow R^1$ of wealth and wants to maximize his expected utility. Then, the agent's payoff function is $u_i(X) = E[g_i(X)]$, $X \in \Gamma$.

REMARK 3.1. In Raiffa (1970) and LaValle (1978), each agent is assumed to have a cardinal utility function of profit, and his objective is to maximize his expected utility. However, some preferences, such as the one in Example 3.2, cannot be represented by a cardinal utility function.

A point $a \in R^N$ is said to be Pareto-inferior to or Pareto-dominated by another point $b \in R^N$, if each component of a is no greater than the corresponding component of b and at least one component of a is less than the corresponding component of b . In other words, we say b is Pareto-superior to a or b Pareto-dominates a . A point is said to be a Pareto-optimal point of a subset of R^N if it is not Pareto-inferior to any other point in the subset. With these concepts, we can now define Pareto-optimality of a sharing rule $\theta(s)$ and an action pair $(s, \theta(s))$.

DEFINITION 3.2. Given an external action s of the supply chain, $\theta^*(s)$ is a Pareto-optimal sharing rule, if

$$(u_1(\Pi_1(s, \theta^*(s))), \dots, u_N(\Pi_N(s, \theta^*(s))))$$

is a Pareto-optimal point of the set

$$\{(u_1(\Pi_1(s, \theta(s))), \dots, u_N(\Pi_N(s, \theta(s))))\}, \theta \in \Theta,$$

where $u_i(\Pi_i(s, \theta(s)))$ is the payoff of the i th agent.

DEFINITION 3.3. $(s^*, \theta^*(s^*))$ is a Pareto-optimal action pair if the agents' payoffs

$$(u_1(\Pi_1(s^*, \theta^*(s^*))), \dots, u_N(\Pi_N(s^*, \theta^*(s^*))))$$

is a Pareto-optimal point of the set

$$\{(u_1(\Pi_1(s, \theta(s))), \dots, u_N(\Pi_N(s, \theta(s))))\}, (s, \theta(s)) \in S \times \Theta.$$

Clearly if $(s^*, \theta^*(s^*))$ is a Pareto-optimal action pair, then $\theta^*(s^*)$ is a Pareto-optimal sharing rule given s^* .

We begin now with an examination of the Pareto-optimal set in a supply chain consisting of risk-neutral agents. If an external action maximizes the supply chain's expected profit, then it is not possible to make one agent get more expected profit without making another agent get less. More specifically, we have the following proposition.

PROPOSITION 3.1. If the agents in a supply chain are all risk neutral, then an action pair $(s, \theta(s))$ is Pareto-optimal

if and only if the channel's external action s maximizes the channel's expected profit.

Proof. The proof follows from the fact that in the risk-neutral case, for each s ,

$$\begin{aligned} \sum u_i(\Pi_i(s, \theta(s))) &= \sum E\Pi_i(s, \theta(s)) \\ &= E \sum \Pi_i(s, \theta(s)) = E\Pi(s). \end{aligned}$$

Thus, every $(s^*, \theta(s^*)) \in S \times \Theta$ is Pareto-optimal provided s^* maximizes $E\Pi(s^*)$. \square

Since agents in a supply chain maximize their respective objectives, the agents' payoffs might not be Pareto-optimal if their objectives are not aligned properly. In this case, it is possible to improve the chain's performance, i.e., achieve Pareto-superior payoffs. The agents can enter into an appropriately designed contract, under which their respective optimizing actions lead to a Pareto-optimal payoff. In the supply chain management literature, a contract is defined to coordinate a supply chain consisting of risk-neutral agents if their respective optimizing external actions under the contract maximize the chain's expected profit. Then, according to Proposition 3.1, a coordinating contract is equivalent to a Pareto-optimal action in the risk-neutral case. It is therefore reasonable to use the notion of Pareto-optimality to define supply chain coordination in the general case.

DEFINITION 3.4. Supply Chain Coordination. A contract agreed upon by the agents of a supply chain is said to *coordinate* the supply chain if the optimizing actions of the agents under the contract:

1. satisfy each agent's reservation payoff constraint, and
2. lead to an action pair $(s^*, \theta^*(s^*))$ that is Pareto-optimal.

Besides Pareto-optimality of a contract, we have introduced the *individual-rationality* or the *participation* constraints as part of the definition of coordination. The constraints ensure that each agent is willing to participate in the contract by requiring that each gets at least his *reservation payoff*. It is clear that each agent's reservation payoff will not be less than his status-quo payoff, which is defined to be his best payoff in the absence of the contract. Thus, we need consider only the subset of Pareto-optimal actions that satisfy these participating constraints. The reservation payoff of an agent plays an important role in bargaining, as we shall see in the next section.

Now we illustrate the introduced concept of coordination by an example.

EXAMPLE 3.4. Consider a supply chain consisting of one supplier and one retailer who faces a newsvendor problem. Before the demand realizes, the supplier decides on his capacity first, and the retailer then prices the product and chooses an order quantity. The sup-

plier and the retailer may enter into a contract that specifies the retailer's committed order quantity and the supplier's refund policy for returned items. In this channel, the external actions are the supplier's capacity selection and the retailer's pricing and ordering decisions. These are denoted as s . The internal actions include decision on the quantity of commitment, the refundable quantity, and the refund credit per item. These internal actions together lead to a sharing rule denoted by $\theta(s)$. Once the contract parameters are determined, the agents in the supply chain choose their respective external actions that maximize their respective payoffs. If $(s, \theta(s))$ satisfies the agents' reservation payoffs and is Pareto-optimal, then the channel is coordinated by the contract.

The definition of coordination proposed here allows agents to have any kind of preference that can be represented by a payoff function satisfying the completeness and transitive axioms. For example, all of the seven kinds of preferences listed in Schweitzer and Cachon (2000), including risk-seeking preferences, are allowed. Since often in practice an agent is either risk neutral or risk averse, we restrict our attention to only these two types.

REMARK 3.2. Our definition applies also to a T -period case. For this, we define the payoff function of player i as

$$u_i(\Pi_i^1(s^*, \theta^*(s^*)), \Pi_i^2(s^*, \theta^*(s^*)), \dots, \Pi_i^T(s^*, \theta^*(s^*))) : \Gamma^T \rightarrow R^1,$$

where $\Pi_i^t(s^*, \theta^*(s^*))$ is agent i 's profit in period t .

4. Coordinating Supply Chains

Each Pareto-optimal action pair $(s, \theta(s))$ results in a vector of payoffs

$$(u_1(\Pi_1(s, \theta(s))), \dots, u_N(\Pi_N(s, \theta(s)))),$$

where $u_i(\Pi_i(s, \theta(s)))$ is the payoff of the i th agent. Let

$$\Psi = \{(u_1(\Pi_1(s, \theta(s))), \dots, u_N(\Pi_N(s, \theta(s)))) | (s, \theta(s))$$

is Pareto-optimal, $(s, \theta(s)) \in S \times \Theta\}$,

denote the set of all Pareto-optimal payoffs, and let $\Phi \subset \Psi$ be the subset of Pareto-optimal payoffs that satisfy all of the participation constraints. We shall refer to Φ as *Pareto-optimal frontier*. We will assume that Φ is not empty.

To coordinate a supply chain, the first step is to obtain the Pareto-optimal frontier Φ . If Φ is not a singleton, then agents bargain to arrive at an element in Φ to which they agree. A coordinating contract is one with a specific set of parameters that achieves the selected solution.

A contract is appealing if it has sufficient flexibility. In Cachon (2002), a coordinating contract is said to be

flexible if the contract, by adjustment of some parameters, allows for any division of the supply chain's expected profit among the risk-neutral agents. This concept can be extended to the general case as follows.

DEFINITION 4.1. A coordinating contract is *flexible* if, by adjustment of some parameters, the contract can lead to any point in Φ .

We shall now develop coordinating contracts in supply chains consisting of two agents: a supplier and a retailer. We shall consider three different cases. In each of these cases, we assume that agents have complete information. In Case 1, the supplier is risk neutral and the retailer has a payoff function in Example 3.2, i.e., the retailer maximizes his expected profit subject to a downside constraint. In Case 2, the supplier and the retailer are both risk averse and each maximizes his own mean-variance trade-off. In Case 3, the supplier and the retailer are both risk averse and each maximizes his own expected concave utility. We consider the first two cases in this section and the third case in Section 5.

In each case, let us denote the retailer's and the supplier's reservation payoffs as $\pi_r \geq 0$ and $\pi_s \geq 0$, respectively. We first obtain Φ and then design a flexible contract that can lead to any point in Φ by adjusting the parameters of the contract.

4.1. Case 1: Risk-Neutral Supplier and Retailer Averse to Downside Risk

We consider the supplier to be risk neutral and the retailer to maximize his expected profit subject to a downside risk constraint. This downside risk constraint requires that the probability of the retailer's profit to be higher than a specified level is not too small.

The risk neutrality assumption on the part of the supplier is reasonable when he is able to diversify his risk by serving a number of independent retailers, which is quite often the case in practice. When the retailers are independent, the supply chain can be divided into a number of sub-chains, each consisting of one supplier and one retailer. This situation, therefore, could be studied as a supply chain consisting of one risk-neutral supplier and one risk-averse retailer.

We say that an action pair $(s, \theta(s))$ is *feasible* if the pair satisfies the retailer's downside risk constraint. We do not need to consider a pair $(s, \theta(s))$ that is not feasible since under the pair the retailer's payoff is $-\infty$ and he would not enter the contract. We denote $\Pi(s)$, $\Pi_r(s, \theta(s))$, and $\Pi_s(s, \theta(s))$ as the profits of the supply chain, the retailer, and the supplier, respectively. Other quantities of interest will be subscripted in the same way throughout the paper, i.e., subscript r will denote the retailer and subscript s will denote the supplier.

Then we have the following result.

THEOREM 4.1. *If the supplier is risk neutral and the retailer maximizes his expected profit subject to a downside risk constraint, then a feasible action pair $(s, \theta(s))$ is Pareto-optimal if and only if the supply chain's expected profit is maximized over the feasible set.*

Proof. ONLY IF: It is sufficient to show that if $E\Pi(s)$ is not maximal over the feasible set, then $(s, \theta(s))$ is not Pareto-optimal.

If $E\Pi(s)$ is not the maximal channel profit, then there exists an s' such that $E\Pi(s') > E\Pi(s)$. Consider the pair $(s', \theta'(s'))$ in which $\Pi_r(s', \theta'(s')) = \Pi_r(s, \theta(s))$ and $\Pi_s(s', \theta'(s')) = \Pi(s') - \Pi_r(s, \theta(s))$, we then get

$$u_r(\Pi_r(s', \theta'(s'))) = E\Pi_r(s, \theta(s)) \text{ and}$$

$$u_s(\Pi_s(s', \theta'(s'))) = E\Pi(s') - E\Pi_r(s, \theta(s)).$$

We can see that

$$u_r(\Pi_r(s', \theta'(s'))) = u_r(\Pi_r(s, \theta(s))) \text{ and}$$

$$u_s(\Pi_s(s', \theta'(s'))) > u_s(\Pi_s(s, \theta(s))).$$

This means that $(s, \theta(s))$ is Pareto-inferior to $(s', \theta'(s'))$, which contradicts with the Pareto-optimality of $(s, \theta(s))$.

IF: Suppose the supply chain's expected profit is maximized. If $(s, \theta(s))$ is not Pareto-optimal, then according to the definition of a Pareto-optimal action pair, there exists a feasible pair $(s'', \theta''(s''))$ that is Pareto-superior to (s, θ) . Since it is Pareto-superior to $(s, \theta(s))$, it is also feasible. Thus,

$$\begin{aligned} u_s(\Pi_s(s'', \theta''(s''))) + u_r(\Pi_r(s'', \theta''(s''))) &= E\Pi(s'') > E\Pi(s) \\ &= u_r(\Pi_r(s, \theta(s))) + u_s(\Pi_s(s, \theta(s))), \end{aligned}$$

which contradicts the fact that $E\Pi(s)$ is the maximal expected channel profit. \square

Let s^* be an action of the channel that maximizes the channel's expected profit, and let $E\Pi_r(s^*)$ be the retailer's payoff. Since the retailer's and the supplier's reservation payoffs are π_r and π_s , respectively, we must impose the participating constraints of the agents on the solutions in Ψ . Thus,

$$E\Pi_r(s^*) \geq \pi_r \text{ and } E\Pi(s^*) - E\Pi_r(s^*) \geq \pi_s. \quad (1)$$

Together with Theorem 4.1, we get

$$\begin{aligned} \Phi &= \{(E\Pi_r(s^*), E\Pi(s^*) - E\Pi_r(s^*)) | E\Pi(s^*) - \pi_s \\ &\geq E\Pi_r(s^*) \geq \pi_r\}. \end{aligned}$$

Clearly, if $E\Pi(s^*) - \pi_s \geq \pi_r$, then Φ is not empty.

In Gan, Sethi, and Yan (2003), we show that a retailer, who is subject to a downside risk constraint, may order a lower quantity from a supplier than that desired by the channel under a wholesale, buy-back, or revenue-sharing contract. Based on an initial contract, a risk-sharing contract is designed, which stipu-

lates the supplier to offer a full refund on unsold items up to a limited quantity. The contract coordinates the supply chain, and requires that both the supplier and the retailer share the risk. Another coordinating contract is possible when $E\Pi_r(s^*)$ exceeds the retailer's target profit α , where s^* is the channel's optimal action. In this case, a contract that provides a payoff of $E\Pi_r(s^*)$ to the retailer and remainder to the supplier coordinates the supply chain. This contract is of two-part tariffs type as defined, for example, in Chopra and Meindl (2001, p. 160). However, if $E\Pi_r(s^*)$ is less than the retailer's target profit α , then the contract does not work since the downside risk constraint of the retailer is not satisfied. But the risk-sharing contract in Gan, Sethi, and Yan (2003) still works, since the retailer's downside risk constraint $P(X \leq \alpha) \leq \beta$ is always satisfied under that contract.

4.2. Case 2: Mean-Variance Suppliers and Retailers

In this case, both the supplier and the retailer maximize their respective mean-variance trade-offs.

Let the retailer's payoff function be

$$E\Pi_r(s, \theta(s)) - \lambda_r V(\Pi_r(s, \theta(s))), \quad (2)$$

and the supplier's payoff function be

$$E\Pi_s(s, \theta(s)) - \lambda_s V(\Pi_s(s, \theta(s))). \quad (3)$$

We first find all Pareto-optimal sharing rules for any given channel's external action s . We show that, regardless of the selected external action s , the optimal sharing rule has the same specific form. Under this form of a sharing rule, we obtain optimal external actions. This procedure results in a Pareto-optimal $(s, \theta(s))$.

We now solve for the Pareto-optimal set for a supply chain consisting of N agents, and then specialize it for supply chains with two agents. We assume that the i th agent's payoff function is

$$E\Pi_i(s, \theta(s)) - \lambda_i V(\Pi_i(s, \theta(s))). \quad (4)$$

To obtain Pareto-optimal sharing rules, we solve

$$\max_{\theta \in \Theta} \left[\sum_i E\Pi_i(s, \theta(s)) - \sum_i \lambda_i V(\Pi_i(s, \theta(s))) \right], \quad (5)$$

s.t.

$$\sum_i \Pi_i(s, \theta(s)) = \Pi(s). \quad (6)$$

The solution of this problem is given in the following proposition.

PROPOSITION 4.1. *A sharing rule θ is a solution of the problem (5)–(6) if and only if*

$$\Pi_i(s, \theta(s)) = \frac{1/\lambda_i}{\sum_j 1/\lambda_j} \Pi(s) + \pi_i, \quad i = 1, \dots, N, \quad (7)$$

almost surely, where π_i , $i = 1 \dots N$, are constants satisfying $\sum_i \pi_i = 0$.

Proof. Because $\sum_i E\Pi_i(s, \theta(s)) = E\Pi(s)$, the problem is equivalent to

$$\min_{\theta \in \Theta} \sum_i \lambda_i V(\Pi_i(s, \theta(s))), \quad (8)$$

s.t.

$$\sum_i \Pi_i(s, \theta(s)) = \Pi(s). \quad (9)$$

It is easy to see that

$$\begin{aligned} \sum_i \lambda_i V(\Pi_i(s, \theta(s))) &= \sum_i \lambda_i V(\Pi_i(s, \theta(s))) \\ &+ \frac{2}{\sum_j 1/\lambda_j} E \left[\Pi(s) \left(\Pi(s) - \sum_i \Pi_i(s, \theta(s)) \right) \right] \\ &= \frac{1}{\sum_j 1/\lambda_j} V(\Pi(s, \theta(s))) \\ &+ \sum_i \lambda_i V \left[\Pi_i(s, \theta(s)) - \frac{1/\lambda_i}{\sum_j 1/\lambda_j} \Pi(s) \right]. \end{aligned} \quad (10)$$

Since the second term on the RHS of (10) is nonnegative, we have shown that

$$\sum_i \lambda_i V(\Pi_i(s, \theta(s))) \geq \frac{1}{\sum_j 1/\lambda_j} V(\Pi(s, \theta(s))) \quad (11)$$

for any feasible $\Pi_i(s, \theta(s))$, $i = 1, \dots, N$. Thus, $(1/\sum_j 1/\lambda_j)V(\Pi(s, \theta(s)))$ provides a lower bound for the objective function (8). Note that a θ satisfies (7) if and only if

$$V \left[\Pi_i(s, \theta(s)) - \frac{1/\lambda_i}{\sum_j 1/\lambda_j} \Pi(s) \right] = 0, \quad i = 1, \dots, N.$$

This means that

$$\sum_i \lambda_i V(\Pi_i(s, \theta(s))) = \frac{1}{\sum_j 1/\lambda_j} V(\Pi(s)) \quad (12)$$

and

$$\sum_i \lambda_i V(\Pi_i(s, \theta(s))) > \frac{1}{\sum_j 1/\lambda_j} V(\Pi(s))$$

for any θ not satisfying (7). \square

For any optimal sharing rule θ^* given in Proposition 4.1, the sum of the agents' payoffs equals

$$E\Pi(s) - \frac{1}{\sum_j 1/\lambda_j} V(\Pi(s)).$$

By adjusting the constants π_i , $i = 1 \dots N$, the sharing

rule θ^* allows for any division of the total payoff among the agents. Therefore, an optimal external action, given θ^* , has to maximize the total payoff, i.e., it must be an action pair of

$$\max_{s \in S} \left[E\Pi(s) - \frac{1}{\sum_j 1/\lambda_j} V(\Pi(s)) \right]. \quad (13)$$

Next, we characterize the set of Pareto-optimal actions by summarizing the results we have gotten.

THEOREM 4.2. *An action pair (s^*, θ^*) is Pareto-optimal if and only if*

$$s^* = \arg \max_{s \in S} \left[E\Pi(s) - \frac{1}{\sum_j 1/\lambda_j} V(\Pi(s)) \right] \quad (14)$$

and, almost surely,

$$\Pi_i(s, \theta^*(s)) = \frac{1/\lambda_i}{\sum_j 1/\lambda_j} \Pi(s) + \pi_i, i = 1, \dots, N. \quad (15)$$

Clearly, if a contract can allocate the channel profit among the N agents proportionally, then the contract along with a side payment scheme can coordinate the supply chain. Moreover, this contract is flexible by adjusting the amounts of side payment.

Theorem 4.3, as a special result of Theorem 4.2, characterizes the set of Pareto-optimal actions for supply chains consisting of one supplier and one retailer.

THEOREM 4.3. *An action pair $(s^*, \theta^*(s^*))$ is Pareto-optimal if and only if*

$$s^* = \arg \max_{s \in S} \left[E\Pi(s) - \frac{\lambda_r \lambda_s}{\lambda_r + \lambda_s} V(\Pi(s)) \right] \quad (16)$$

and, almost surely,

$$\Pi_r(s^*, \theta^*(s^*)) = \frac{\lambda_s}{\lambda_r + \lambda_s} \Pi(s^*) + \pi_0, \quad (17)$$

$$\Pi_s(s^*, \theta^*(s^*)) = \frac{\lambda_r}{\lambda_r + \lambda_s} \Pi(s^*) - \pi_0, \quad (18)$$

where π_0 is a side payment from the supplier to the retailer.

It follows from Theorem 4.3 that, under any Pareto-optimal solution, the retailer gets a fixed proportion $\lambda_s/(\lambda_r + \lambda_s)$ of the channel profit plus π_0 and the supplier gets the remaining profit, i.e., $\lambda_r/(\lambda_r + \lambda_s)$ of the channel profit minus π_0 . If $\lambda_r > \lambda_s$, i.e., the retailer is more risk-averse than the supplier, then the supplier takes a greater proportion of the channel profit. In other words, the agent with a lower risk aversion takes a higher proportion of the total channel profit than the other one does. The side payment, which is determined by the respective bargaining powers of the agents, determines the agents' final payoffs.

According to Theorem 4.3,

$$\Psi = \left\{ \left(\frac{\lambda_s}{\lambda_r + \lambda_s} u(s^*) + \pi_0, \frac{\lambda_r}{\lambda_r + \lambda_s} u(s^*) - \pi_0 \right) \mid \pi_0 \in \mathbb{R}^1 \right\},$$

where $u(s^*)$ represents

$$E\Pi(s^*) - \lambda V(\Pi(s^*)).$$

Since the retailer's and the supplier's reservation payoffs are π_r and π_s , respectively, π_0 has to satisfy the participating constraints of the agents. Thus,

$$\frac{\lambda_s}{\lambda_r + \lambda_s} u(s^*) + \pi_0 \geq \pi_r \text{ and } \frac{\lambda_r}{\lambda_r + \lambda_s} u(s^*) - \pi_0 \geq \pi_s.$$

Then Φ can be represented by the set

$$\left\{ \left(\frac{\lambda_s}{\lambda_r + \lambda_s} u(s^*) + \pi_0, \frac{\lambda_r}{\lambda_r + \lambda_s} u(s^*) - \pi_0 \right) \text{ such that } \frac{\lambda_r}{\lambda_r + \lambda_s} u(s^*) - \pi_s \geq \pi_0 \geq \pi_r - \frac{\lambda_s}{\lambda_r + \lambda_s} u(s^*) \right\}.$$

Furthermore, if

$$\frac{\lambda_r}{\lambda_r + \lambda_s} u(s^*) - \pi_s \geq \pi_r - \frac{\lambda_s}{\lambda_r + \lambda_s} u(s^*),$$

i.e., if $\pi_r + \pi_s \leq u(s^*)$, then Φ is not empty.

The problem considered thus far is quite general, in the sense that the external action s is rather an abstract one that can include such decisions as order quantity, item price, etc. We next consider a special case. Here the retailer faces a newsvendor problem and makes a single purchase order of a product from the supplier at the beginning of a period, who in turn produces and delivers the order to the retailer before the selling season commences. Let p denote the price per unit, c the supplier's production cost, v the salvage value, and q the retailer's order quantity. In this problem, the supply chain's external action is the retailer's order quantity q .

According to Theorem 4.3, the coordinating contract should allocate the profit in the same proportion for every realization of the channel profit in the absence of any side payment. We shall call such a sharing rule a *proportional sharing rule*. Here we only examine buy-back and revenue-sharing contracts. With a buy-back contract, the supplier charges the retailer a wholesale price per unit, but he pays the retailer a credit for every unsold unit at the end of the season. With a revenue-sharing contract, the supplier charges a wholesale price per unit purchased, and the retailer gives the supplier a percentage of his revenue. See Pasternack (1985) and Cachon and Lariviere (2002) for details on these contracts. In the following, we see that both buy-back and revenue-sharing contracts allocate the channel profit proportionally.

PROPOSITION 4.2. *If the parameters of a revenue-sharing contract satisfy*

$$w = \phi c, \quad (19)$$

then the retailer's fractional share of the supply chain profit is ϕ and the supplier's share is $1 - \phi$.

Proof. Let D denote the demand faced by the retailer. Then the supply chain's profit is

$$\begin{cases} pD + (q - D)v - cq & \text{if } D \leq q, \\ pq - cq, & \text{if } D > q. \end{cases} \quad (20)$$

On the other hand, the retailer's profit is

$$\begin{cases} \phi pD + \phi(q - D)v - wq & \text{if } D \leq q, \\ \phi pq - wq, & \text{if } D > q. \end{cases} \quad (21)$$

By using $w = \phi c$ into (21), we can see that the retailer gets the proportion ϕ of the supply chain's profit for every realization of the demand. \square

Cachon and Lariviere (2002) prove that, for each coordinating revenue-sharing contract, there exists a unique buy-back contract that provides the same profit as in the revenue-sharing contract for every demand realization. They show that the buy-back contract's parameters have the form

$$b = (1 - \phi)(p - v), \quad (22)$$

$$w = p(1 - \phi) + \phi c, \quad (23)$$

where b is the refund to the retailer for each unsold unit, and ϕ is the retailer's share of the channel profit in the revenue-sharing contract. It is easy to see that the same result holds here as well.

PROPOSITION 4.3. *A buy-back contract allocates the channel profit (a random variable) proportionally. If the contract parameters satisfy (22) and (23), then the retailer's fractional share is ϕ and the supplier's is $1 - \phi$ under this contract.*

Nagarajan and Bassok (2003) obtain the Nash Bargaining solution in the risk-neutral case. According to their results, if both the retailer and the supplier are risk neutral, the retailer's share of the supply chain profit is

$$\phi \Pi(s) = [\Pi(s) - \pi_s + \pi_r] / 2. \quad (24)$$

Under a buy-back or a revenue-sharing contract, the retailer's problem is

$$\max_{s \in S} \left[\pi_0 + \frac{\lambda_s}{\lambda_r + \lambda_s} [E\Pi(s) - \lambda V(\Pi(s))] \right]. \quad (25)$$

For a given π_0 , the retailer's problem (25) becomes

$$\max_{s \in S} [E\Pi(s) - \lambda V(\Pi(s))], \quad (26)$$

which has been solved by Lau (1980) and Chen and Federgruen (2000). Since the solution is s^* , the retailer would choose the optimal external action voluntarily. So, we can state the following two theorems.

THEOREM 4.4. *If the parameters of a revenue-sharing contract satisfy*

$$w = \frac{\lambda_s}{\lambda_r + \lambda_s} c, \quad (27)$$

then the revenue-sharing contract along with a side payment π_0 to the retailer coordinates the supply chain. The profit allocation is given in (17)–(18).

THEOREM 4.5. *If the parameters of a buy-back contract satisfy*

$$b = \frac{\lambda_r}{\lambda_r + \lambda_s} (p - v), \quad (28)$$

$$w = p \frac{\lambda_r}{\lambda_r + \lambda_s} + \frac{\lambda_s}{\lambda_r + \lambda_s} c, \quad (29)$$

then the buy-back contract along with a side payment π_0 to the retailer coordinates the supply chain. The profit allocation is given in (17)–(18).

If we compare (19) to (27) and (22)–(23) to (28)–(29), we see that, while the fraction ϕ in the risk-neutral case can vary as long as the participation constraints are satisfied, the fractional share the retailer receives in the mean-variance case is fixed at $\lambda_s / (\lambda_r + \lambda_s)$; and it is this fraction that is used in the formulations (27)–(29) for the parameters of the coordinating contracts. Thus, in the mean-variance case, the participation constraints of the agents as well as the flexibility of the contracts are achieved only via the side payment π_0 .

The contracts obtained in Theorems 4.5 and 4.4, when $\lambda_r = \lambda_s = 0$, reduce to the standard contracts obtained in the risk-neutral case, because the fraction $\lambda_r / (\lambda_r + \lambda_s)$ can take any value in $[0, 1]$. In particular, if the supplier is risk neutral and the retailer is risk averse, i.e., $\lambda_s = 0$, the fraction $\lambda_r / (\lambda_r + \lambda_s) = 1$, which means that the supplier takes the entire channel profit and gives a side payment to the retailer. In this case, it is Pareto-optimal for the supplier to bear all of the risk. Since the retailer's profit is a side payment from the supplier, the supplier's expected profit is the channel's profit minus that payment. Therefore, the supplier's payoff is maximized when the channel's expected profit is maximized. Thus, we have a coordinating contract under which the supplier and the retailer execute s^* , the retailer gets a constant profit π_0 , and the supplier gets the remaining profit.

In the following example, we design a coordinating contract according to Theorem 4.5. We also obtain the optimal ordering quantity and determine the required side payment.

EXAMPLE 4.1. Consider a supply chain consisting of one retailer and one supplier. The retailer faces a newsvendor problem and makes a single purchase order of a product from the supplier at the beginning of a period, who in turn produces and delivers the order to the retailer before the selling season. Suppose that the demand D is uniformly distributed on some interval, which without loss of generality, can be taken as interval $[0, 1]$. Thus, the distribution function $F(x) = x$ for $0 \leq x \leq 1$ and $F(x) = 1$ for $x \geq 1$. We confine the ordering quantity q to be in $[0, 1]$. Let the unit price p be 100, the supplier's production cost c be 60, and the salvage value v be 20. Let the retailer's and the supplier's payoff functions be, respectively,

$$E\Pi_r - \lambda_r V(\Pi_r) \text{ and } E\Pi_s - \lambda_s V(\Pi_s), \quad (30)$$

where $\lambda_r = 0.05$ and $\lambda_s = 0.01$. We assume that the agents have equal bargaining powers in the sense that their payoffs are equal.

According to Theorem 4.3, the retailer's payoff is

$$\frac{\lambda_s}{\lambda_r + \lambda_s} \left[E\Pi(q) - \frac{\lambda_r \lambda_s}{\lambda_r + \lambda_s} V(\Pi(q)) \right] + \pi_0, \quad (31)$$

where $\Pi(q)$ is the channel's profit when the retailer's ordering quantity is q , $0 \leq q \leq 1$. Thus, the retailer's optimal order quantity is a q^* that maximizes (31). From Chen and Federgruen (2000), we have

$$E\Pi(q) = 40q - 80q^2 \text{ and } V(\Pi(q)) = 6400(q^3/3 - q^4/4). \quad (32)$$

With this, the retailer's problem is

$$\max_{0 \leq q \leq 1} \left[40q - 80q^2 - \frac{160}{9} q^3 + \frac{40}{3} q^4 \right], \quad (33)$$

and $q^* = 0.236$. According to Theorem 4.5, the retailer's and the supplier's payoffs are $0.799 + \pi_0$ and $3.99 - \pi_0$, respectively. It is easy to see that $\pi_0 = 1.598$ equalizes their payoffs, as has been assumed.

5. Coordinating a Supply Chain Consisting of Agents with Concave Utility Functions

In this case, we assume that agent i has an increasing concave utility function $g_i(\cdot)$ of his profit, and wants to maximize his expected utility, $i = r, s$. Then his payoff function is $E[g_i(\cdot)]$.

To compute the set of Pareto-optimal actions, we first find the Pareto-optimal sharing rules given an external action s . According to the group decision theory literature (Wilson 1968 and Raiffa 1970), the problem can be formulated as follows:

$$\max_{\theta \in \Theta} [a_r E g_r(\Pi_r(s, \theta(s))) + a_s E g_s(\Pi_s(s, \theta(s)))], \quad (34)$$

s. t.

$$\Pi_r(s, \theta(s)) + \Pi_s(s, \theta(s)) = \Pi(s), \quad (35)$$

where $a_r, a_s > 0$, $a_r + a_s = 1$. The specification of (a_r, a_s) is derived from their respective bargaining powers. By varying a_r and a_s we can get all possible Pareto-optimal sharing rules Ψ_s denoted as the set

$$\{(u_r(\Pi_r(s, \theta(s))), u_s(\Pi_s(s, \theta(s)))) \text{ such that } \theta \text{ is Pareto-optimal, } \theta \in \Theta\}. \quad (36)$$

Clearly, each point in Ψ_s represents, given s , the agents' payoffs under a Pareto-optimal sharing rule. Then we can get Ψ , which is the set of Pareto-optimal points of the set $\cup_{s \in S} \Psi_s$. According to Definition 3.3, any action pair that leads to a point in Ψ is Pareto-optimal.

It is well known that the problem of maximizing the expected quadratic utility can be reduced to one of maximizing a mean-variance trade-off. Therefore, when both agents' utility functions are quadratic, we can coordinate the channel with the contracts developed in Section 4.2. Levy and Markowitz (1979) show that a utility function exhibiting constant risk aversion, particularly of the form $\log x$ or x^α , $0 < \alpha < 1$, can be approximated by a quadratic function. Thus, if both agents' utility functions are of this type, we can use the contracts in Section 4.2 to coordinate the channel approximately.

For general utility functions $g_i(\cdot)$, $i = r, s$, it is not always possible to get Pareto-sharing rules in closed form. Indeed, in the literature, closed-form solutions exist only in a few special cases. In the next subsection, we specialize the model by assuming the case of exponential utility functions, and obtain an explicit coordinating contract. We also obtain the Nash Bargaining solution. Finally, we design a flexible contract to coordinate the supply chain.

5.1. Characterizing the Pareto-Optimal Set

Let the retailer's and the supplier's utility functions be, respectively,

$$g_r(x) = 1 - e^{-\lambda_r x} \text{ and } g_s(x) = 1 - e^{-\lambda_s x}. \quad (37)$$

Then the absolute risk aversion measure for the retailer and the supplier are λ_r and λ_s , respectively; see Pratt (1964). We want to find a Pareto-optimal sharing rule for any given channel's external action s . Raiffa (1970) solves the problem (34)–(35), which implies the following result.

PROPOSITION 5.1. For a given external action s for the channel under consideration, a sharing rule $\theta^*(s)$ is a Pareto-optimal sharing rule if and only if

$$\Pi_r(s, \theta^*(s)) = \frac{\lambda_s}{\lambda_r + \lambda_s} \Pi(s) - \lambda \ln \frac{a_s \lambda_s}{a_r \lambda_r} \quad (38)$$

and

$$\Pi_s(s, \theta^*(s)) = \frac{\lambda_r}{\lambda_r + \lambda_s} \Pi(s) + \lambda \ln \frac{a_s \lambda_s}{a_r \lambda_r}, \quad (39)$$

almost surely, where $a_r, a_s > 0$, $a_r + a_s = 1$, and $\lambda = 1/(\lambda_r + \lambda_s)$.

Thus, we can get the set

$$\Psi_s = \{(u_r(\Pi_r(s, \theta^*(s))), u_s(\Pi_s(s, \theta^*(s)))) \text{ such that } a_r, a_s > 0, a_r + a_s = 1\},$$

where

$$u_r(\Pi_r(s, \theta^*(s))) = 1 - \exp\left(\frac{\lambda_r}{\lambda_r + \lambda_s} \ln \frac{a_s \lambda_s}{a_r \lambda_r}\right) E \exp\left(-\frac{\lambda_r \lambda_s \Pi(s)}{\lambda_r + \lambda_s}\right), \quad (40)$$

$$u_s(\Pi_s(s, \theta^*(s))) = 1 - \exp\left(\frac{\lambda_s}{\lambda_r + \lambda_s} \ln \frac{a_r \lambda_r}{a_s \lambda_s}\right) E \exp\left(-\frac{\lambda_r \lambda_s \Pi(s)}{\lambda_r + \lambda_s}\right). \quad (41)$$

Since both the retailer and the supplier's payoff functions decrease with

$$E \exp\left(-\frac{\lambda_r \lambda_s \Pi(s)}{\lambda_r + \lambda_s}\right),$$

it is easy to check that

$$\Psi = \{(u_r(\Pi_r(s^*, \theta^*(s^*))), u_s(\Pi_s(s, \theta^*(s)))) | a_r, a_s > 0, a_r + a_s = 1\}, \quad (42)$$

where s^* is the solution of the problem

$$\min_{s \in S} E \exp\left(-\frac{\lambda_r \lambda_s \Pi(s)}{\lambda_r + \lambda_s}\right). \quad (43)$$

Now the supply chain's external problem has been transformed to problem (43). This problem has been studied in the literature in some special situations. Bouakiz and Sobel (1992) have shown that a base-stock policy is optimal in a multi-period newsvendor problem, when the newsvendor has an exponential utility function. Eeckhoudt, Gollier, and Schlesinger (1995) discuss the situation in which the entity faces a newsvendor problem, and they prove that the newsvendor orders less than that in the risk-neutral case. Agrawal and Seshadri (2000a) consider the entity's price and inventory decisions jointly in a newsvendor framework.

REMARK 5.1. *Although we have got proportional sharing rules for the above case and the second case in Section 4, the Pareto-optimal sharing rules usually are not proportional for any two utility functions (Raiffa 1970). Moreover,*

the Pareto-sharing rules may depend on the channel's external action. See Wilson (1968), Raiffa (1970), and LaValle (1978) for further details on Pareto-optimal sharing rules.

Now, we summarize the results in the following theorem.

THEOREM 5.1. *An action pair $(s^*, \theta^*(s^*))$ is Pareto-optimal if and only if*

$$s^* = \arg \min_{s \in S} E \exp\left(-\frac{\lambda_r \lambda_s \Pi(s)}{\lambda_r + \lambda_s}\right) \quad (44)$$

and, almost surely,

$$\Pi_r(s^*, \theta^*(s^*)) = \frac{\lambda_s}{\lambda_r + \lambda_s} \Pi(s^*) - \lambda \ln \frac{a_s \lambda_s}{a_r \lambda_r}, \quad (45)$$

$$\Pi_s(s^*, \theta^*(s^*)) = \frac{\lambda_r}{\lambda_r + \lambda_s} \Pi(s^*) + \lambda \ln \frac{a_s \lambda_s}{a_r \lambda_r}, \quad (46)$$

where $a_r, a_s > 0$, $a_r + a_s = 1$.

It follows from Theorem 5.1 that, under any Pareto-optimal solution, the retailer and the supplier get fixed proportions of the channel profit minus/plus a side payment. The side payment depends on their risk-aversion measures and their bargaining powers. In particular, if $a_s \lambda_s = a_r \lambda_r$, then there is no side payment. Moreover, in the absence of a side payment, if $\lambda_r > \lambda_s$, i.e., if the retailer is more risk averse than the supplier, then the supplier takes a greater proportion of the channel profit.

5.2. Bargaining Issue

We have gotten Pareto-optimal payoffs set Ψ in (42). Since $u_r < 1$ and $u_s < 1$ from (40) and (41), we assume that $\pi_r < 1$ and $\pi_s < 1$. The participating constraints of the agents are

$$u_r(\Pi_r(s^*, \theta^*(s^*))) \geq \pi_r, \text{ and } u_s(\Pi_s(s^*, \theta^*(s^*))) \geq \pi_s. \quad (47)$$

Conditions (47) are equivalent to

$$(1 - \pi_r)/E \exp\left(-\frac{\lambda_r \lambda_s \Pi(s^*)}{\lambda_r + \lambda_s}\right) \geq \left(\frac{a_s \lambda_s}{a_r \lambda_r}\right)^{\lambda_r/(\lambda_r + \lambda_s)} \geq E \exp\left(-\frac{\lambda_r \lambda_s \Pi(s^*)}{\lambda_r + \lambda_s}\right)/(1 - \pi_s). \quad (48)$$

Then Φ can be represented by

$$\{(u_r(\Pi_r(s^*, \theta^*(s^*))), u_s(\Pi_s(s^*, \theta^*(s^*)))) | a_r, a_s > 0, a_r + a_s = 1, \text{ and (48) is satisfied.}\}. \quad (49)$$

If $(1 - \pi_r)/E \exp[-\lambda_r \lambda_s \Pi(s^*)/(\lambda_r + \lambda_s)] \geq E \exp[-\lambda_r \lambda_s \Pi(s^*)/(\lambda_r + \lambda_s)]/(1 - \pi_s)$, i.e., if

$$(1 - \pi_r)(1 - \pi_s) \geq \left\{ E \exp\left(-\frac{\lambda_r \lambda_s \Pi(s^*)}{\lambda_r + \lambda_s}\right) \right\}^2,$$

then Φ is not empty.

Nagarajan and Bassok (2003) use the Nash Bargaining concept to deal with bargaining issue in the risk-neutral case. Here, we use the same concept to deal with the bargaining issue in a risk-averse case. According to Nash (1950), a bargaining solution satisfies the following eight axioms.

1. An agent offered two possible anticipations can decide which is preferable or that they are equally desirable.

2. The ordering thus produced is transitive; if A is better than B and B is better than C , then A is better than C .

3. Any probability combination of equally desirable states is just as desirable as either.

4. If A , B , and C are as in Axiom 2, then there is a probability combination of A and C which is just as desirable as B . This amounts to an assumption of continuity.

5. If $0 \leq p \leq 1$ and A and B are equally desirable, then $pA + (1 - p)C$ and $pB + (1 - p)C$ are equally desirable. Also, A may be substituted for B in any desirability ordering relationship satisfied by B .

6. Let S , which is compact and convex and includes the point (π_r, π_s) , be the agents' payoffs set and let $c(S)$ denote the bargaining solution point in this set. If α is a point in S such that there exists another point β in S with the property $u_r(\beta) > u_r(\alpha)$ and $u_s(\beta) > u_s(\alpha)$, then $\alpha \notin c(S)$.

7. If the set T contains the set S and $c(T)$ is in S , then $c(T) = c(S)$.

8. If S is symmetric, i.e., its graph is symmetrical with respect to the line $u_r = u_s$, then $c(S)$ is a point on the line $u_r = u_s$.

Clearly, exponential utilities in (37) satisfy the first five axioms. The agents' payoff set S is

$$\{(u_r(\Pi_r(s, \theta(s))), u_s(\Pi_s(s, \theta(s)))) | (s, \theta(s)) \in S \times \Theta\}.$$

We assume that $\Pi(s)$ be continuous in s , so that S is compact. Now we prove the convexity of S by showing that its frontier is a concave curve. The Pareto-frontier of this set is given in (49), where

$$u_r(\Pi_r(s^*, \theta^*(s^*))) = 1 - \exp\left(\frac{\lambda_r}{\lambda_r + \lambda_s} \ln \frac{a_s \lambda_s}{a_r \lambda_r}\right) E \exp\left(-\frac{\lambda_r \lambda_s \Pi(s^*)}{\lambda_r + \lambda_s}\right), \quad (50)$$

$$u_s(\Pi_s(s^*, \theta^*(s^*))) = 1 - \exp\left(\frac{\lambda_s}{\lambda_r + \lambda_s} \ln \frac{a_r \lambda_r}{a_s \lambda_s}\right) E \exp\left(-\frac{\lambda_r \lambda_s \Pi(s^*)}{\lambda_r + \lambda_s}\right). \quad (51)$$

From (50) and (51),

$$(1 - u_r)^{(\lambda_r + \lambda_s)/\lambda_r} (1 - u_s)^{(\lambda_r + \lambda_s)/\lambda_s} = \left\{ E \exp\left(-\frac{\lambda_r \lambda_s \Pi(s^*)}{\lambda_r + \lambda_s}\right) \right\}^{(\lambda_r + \lambda_s)^2/\lambda_r \lambda_s}. \quad (52)$$

Clearly, the curve represented by (52) is concave since the right side is a constant.

Axiom 6 assures that the solution is Pareto-optimal, Axiom 8 expresses the bargaining equality of bargaining skills. Nash (1950) shows that the bargaining solution point is the point that is the solution of the problem

$$\max_{(u_r, u_s) \in S} (u_r - \pi_r)(u_s - \pi_s). \quad (53)$$

Since the solution has to be Pareto-optimal, (53) is equivalent to

$$\max_{(u_r, u_s) \in \Phi} (u_r - \pi_r)(u_s - \pi_s). \quad (54)$$

Now we solve the problem (54). Let $E(u)$ represent $E \exp[-\lambda_r \lambda_s \Pi(s^*)/(\lambda_r + \lambda_s)]$. Then,

$$\begin{aligned} & (u_r - \pi_r)(u_s - \pi_s) \\ &= (1 - \pi_r)(1 - \pi_s) + \exp\left(\frac{\lambda_r - \lambda_s}{\lambda_r + \lambda_s}\right) [E(u)]^2 \\ & - E(u) \left[(1 - \pi_r) \left(\frac{a_s \lambda_s}{a_r \lambda_r}\right)^{-\lambda_s/(\lambda_r + \lambda_s)} + (1 - \pi_s) \left(\frac{a_s \lambda_s}{a_r \lambda_r}\right)^{\lambda_r/(\lambda_r + \lambda_s)} \right]. \end{aligned}$$

Thus, problem (54) is equivalent to

$$\min_{a_r, a_s} \left[(1 - \pi_r) \left(\frac{a_s \lambda_s}{a_r \lambda_r}\right)^{-\lambda_s/(\lambda_r + \lambda_s)} + (1 - \pi_s) \left(\frac{a_s \lambda_s}{a_r \lambda_r}\right)^{\lambda_r/(\lambda_r + \lambda_s)} \right]. \quad (55)$$

Thus, when

$$\frac{a_s \lambda_s}{a_r \lambda_r} = \frac{(1 - \pi_r) \lambda_s}{(1 - \pi_s) \lambda_r}, \quad (56)$$

$(u_r - \pi_r)(u_s - \pi_s)$ is maximized. Therefore, the Nash Bargaining solution is represented by the payoffs given in (50) and (51) with parameters (a_r, a_s) that satisfy $a_r > 0$, $a_s > 0$, $a_r + a_s = 1$, and (56).

From (56), we can see that $a_s \lambda_s / a_r \lambda_r$ decreases in the retailer's reservation payoff π_r and increases in the supplier's reservation payoff π_s . Thus, the retailer's payoff increases in π_r and decreases in π_s . The same property holds for the supplier. These properties imply that, in the Nash Bargaining solution, each agent's payoff increases with his own reservation payoff and decreases with the other agent's reservation payoff. When $(1 - \pi_r) \lambda_s = (1 - \pi_s) \lambda_r$, then $\lambda \ln(a_s \lambda_s / a_r \lambda_r) = 0$, i.e. the side payment disappears. Finally, we can see that the Nash bargaining solution assigns a player a higher payoff when the other player becomes more risk averse.

5.3. Designing a Coordinating Contract

For the special supply chain considered in Section 4.2, we can use either a buy-back or a revenue-sharing contract to allocate the channel profit. Under either of these contracts, the retailer's problem is:

$$\max_{s \in S} \left[1 - E \exp\left(-\frac{\lambda_r \lambda_s \Pi(s)}{\lambda_r + \lambda_s}\right) \right].$$

This problem is equivalent to problem (43), which implies that the retailer would voluntarily choose the optimal external action s^* . So we have the following two results.

THEOREM 5.2. *If the parameters of a buy-back contract satisfy*

$$b = \frac{\lambda_r}{\lambda_r + \lambda_s} (p - v),$$

$$w = p \frac{\lambda_r}{\lambda_r + \lambda_s} + \frac{\lambda_s}{\lambda_r + \lambda_s} c,$$

then the buy-back contract along with the side payment $\lambda \ln(a_s \lambda_s / a_r \lambda_r)$ to the supplier coordinates the supply chain. The profit allocation is given in (45)–(46).

THEOREM 5.3. *If the parameters of a revenue-sharing contract satisfy*

$$w = \frac{\lambda_s}{\lambda_r + \lambda_s} c,$$

then the revenue-sharing contract along with the side payment $\lambda \ln(a_s \lambda_s / a_r \lambda_r)$ to the supplier coordinates the supply chain. The profit allocation is given in (45)–(46).

Note that, if a_r and a_s satisfy condition (56), then both the revenue-sharing and the buy-back contracts achieve the Nash Bargaining solution. By adjusting the bargaining coefficients a_r and a_s , one can attain any point in Φ . Thus, both these contracts are flexible.

We should note that, in general, Pareto-optimal sharing rules are not proportional as in the case with exponential utility functions. Wilson (1968) provides a necessary and sufficient condition for Pareto-optimality of a sharing rule in a channel with N agents. The condition is stated in the following theorem.

THEOREM 5.4. *Given an external action s , a necessary and sufficient condition for Pareto-optimality of a sharing rule is that there exists nonnegative weights $\mathbf{a} = (a_1, a_2, \dots, a_N)$ and a function $\mu : R^1 \rightarrow R^1$, such that*

$$\sum_i \Pi_i(s, \theta(s)) = \Pi(s), \tag{57}$$

almost surely, and for each i

$$a_i g'_i(\Pi_i(s, \theta(s))) = \mu(\Pi(s)), \tag{58}$$

almost surely.

In what follows, we give an example in which the sharing rule is not of the form of (45)–(46). Here we see that the Pareto-optimal sharing rule depends on the

realized channel profit, i.e., it depends on the chosen external action as well as the realized random event.

EXAMPLE 5.1. Let

$$g_r(x) = 1 - e^{-\lambda_r x}, \tag{59}$$

$$g_s(x) = \begin{cases} 1 - e^{-\lambda_s x}, & x \leq x_0, \\ 1 - e^{-x_0 \lambda_s} + \lambda_s e^{-x_0 \lambda_s} (x - x_0), & x > x_0. \end{cases} \tag{60}$$

In this example, the retailer’s utility function is the same as in (37), but the supplier’s utility is changed in a way that his risk attitude is the same as in (37) at low profit levels and he is risk neutral at higher profit levels.

PROPOSITION 5.2. *For a given external action s for the channel under consideration, a sharing rule θ^* is a Pareto-optimal sharing rule if and only if*

$$\Pi_r(s, \theta^*(s)) = \begin{cases} \frac{\lambda_s}{\lambda_r + \lambda_s} \Pi(s) - \lambda \ln \frac{a_s \lambda_s}{a_r \lambda_r}, & \Pi(s) \leq \frac{\lambda_r + \lambda_s}{\lambda_r} \left(x_0 - \lambda \ln \frac{a_s \lambda_s}{a_r \lambda_r} \right) \\ \frac{\lambda_s}{\lambda_r} x_0 - \lambda \ln \frac{a_s \lambda_s}{a_r \lambda_r}, & \Pi(s) > \frac{\lambda_r + \lambda_s}{\lambda_r} \left(x_0 - \lambda \ln \frac{a_s \lambda_s}{a_r \lambda_r} \right) \end{cases}, \tag{61}$$

$$\Pi_s(s, \theta^*(s)) = \begin{cases} \frac{\lambda_r}{\lambda_r + \lambda_s} \Pi(s) + \lambda \ln \frac{a_s \lambda_s}{a_r \lambda_r}, & \Pi(s) \leq \frac{\lambda_r + \lambda_s}{\lambda_r} \left(x_0 - \lambda \ln \frac{a_s \lambda_s}{a_r \lambda_r} \right) \\ \Pi(s) - \frac{\lambda_s}{\lambda_r} x_0 + \lambda \ln \frac{a_s \lambda_s}{a_r \lambda_r}, & \Pi(s) > \frac{\lambda_r + \lambda_s}{\lambda_r} \left(x_0 - \lambda \ln \frac{a_s \lambda_s}{a_r \lambda_r} \right) \end{cases}, \tag{62}$$

almost surely, where $a_r, a_s > 0$, $a_r + a_s = 1$, and $\lambda = 1/(\lambda_r + \lambda_s)$.

Proof. Let

$$\mu(t) = \begin{cases} (a_r \lambda_r)^{\lambda \lambda_s} (a_s \lambda_s)^{\lambda \lambda_r} \exp\left(\frac{\lambda_r \lambda_s t}{\lambda_r + \lambda_s}\right), & t \leq \frac{\lambda_r + \lambda_s}{\lambda_r} \left(x_0 - \lambda \ln \frac{a_s \lambda_s}{a_r \lambda_r} \right), \\ (a_r \lambda_r)^{\lambda \lambda_s} (a_s \lambda_s)^{\lambda \lambda_r} \exp\left[\lambda_s \left(x_0 - \lambda \ln \frac{a_s \lambda_s}{a_r \lambda_r} \right)\right], & t > \frac{\lambda_r + \lambda_s}{\lambda_r} \left(x_0 - \lambda \ln \frac{a_s \lambda_s}{a_r \lambda_r} \right). \end{cases}$$

Then, according to Theorem 5.4, (61) and (62) are Pareto-optimal since conditions (57) and (58) are satisfied. \square

We can see that, under the Pareto-optimal sharing rule, the retailer’s profit increases linearly with the channel’s realized profit when the latter is below a certain level, and remains unchanged thereafter. This

is not a proportional sharing rule, and consequently, neither a buy-back nor a revenue-sharing contract along with side payments would coordinate the channel. It appears that new contract forms need to be designed to achieve coordination in such cases.

In order to obtain s^* , we outline the following procedure. First, we compute Ψ_s for each $s \in S$ according to (36), and then we find the Pareto-optimal frontier of the set $\cup_{s \in S} \Psi_s$. Any action pair $(s^*, \theta^*(s^*))$ that leads to a point on this frontier is Pareto-optimal. Note that s^* may not be unique.

To illustrate this procedure, let us assume $S = \{s_1, s_2, s_3, s_4, s_5\}$ for convenience in exposition. Suppose that the sets Ψ_s for $s \in S$ are as shown in Figure 1. Then the frontier consisting of Pareto-optimal solutions is shown as the bold-faced boundary in the figure. Construction of such a frontier in general would require development of numerical procedures. This is not the focus of this paper, and it is a topic for future research.

6. Discussion

One of our main findings is that, in any Pareto-optimal joint action, the retailer and the supplier must share the risk appropriately. Specifically, the less risk averse an agent is, the more risk he assumes by taking a larger portion of the channel's random profit. The agents' final payoffs can be adjusted by a side payment depending on their respective bargaining powers. In the extreme case when one of the agents is risk neutral, then that agent may assume all of the risk.

Owing to the risk-sharing effect, the supply chain, when considered as a single entity, is less risk averse than either the risk-averse retailer or the risk-averse supplier if considered as the single owner of the whole channel. For example, in Case 2, the channel's problem according to (25) is equivalent to solving the problem

$$\max_{s \in S} [E(\Pi(s)) - \lambda V(\Pi(s))].$$

If the retailer or the supplier were to own the channel, he would solve the problem

$$\max_{s \in S} [E\Pi(s) - \lambda_r V(\Pi(s))] \quad \text{or} \quad \max_{s \in S} [E\Pi(s) - \lambda_s V(\Pi(s))].$$

Since $\lambda < \lambda_r$ and $\lambda < \lambda_s$, the channel is less risk averse than either the retailer or the supplier. In the extreme case when the supplier is risk neutral, the channel's Pareto-optimal external action, as in the risk-neutral case, always maximizes the channel's expected profit. This result can be generalized to the more general case in which there exists a risk-neutral agent in a supply chain consisting of N agents. Under an external action that maximizes the channel's expected profit, there exists a Pareto-optimal sharing rule that lets the risk-neutral agent take all the channel profit and give each of the other agents a side payment. The amount of the side payment an agent receives depends on his bargaining power.

Although we have obtained explicit solutions for each of the cases in Sections 4 and 5, it is not always possible to do so. For example, if the retailer maximizes his mean-variance trade-off and the supplier maximizes his expected exponential utility, then it does not seem possible to find a closed-form solution. Nevertheless, one can see that the sharing of the risk has to be invoked to get a Pareto-optimal solution.

REMARK 6.1. *An alternative way to consider coordination is to first construct a supply chain's utility function by assigning weights to individuals' utility functions. Thus,*

$$u(\Pi(s)) = \sum_{i=1}^N \alpha_i u_i(\Pi(s)), \quad \alpha_i \in [0, 1], \quad \sum_{i=1}^N \alpha_i = 1,$$

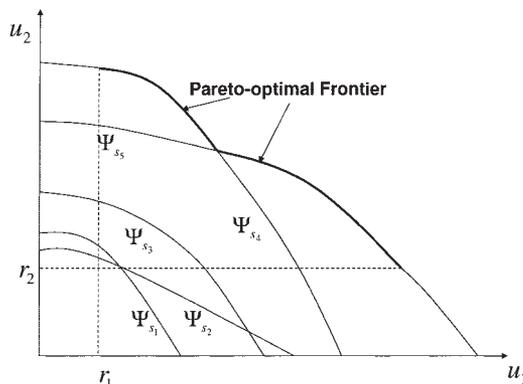
where the agent i 's utility function is denoted as $u_i(\Pi(s))$, $i = 1, 2, \dots, N$. In this case, each agent's identity is not preserved and the critical issue is how to determine the weights α_i , $i = 1, 2, \dots, N$. Once the optimal action for the supply chain is obtained by maximizing $u(\Pi(s))$, the profit could be allocated according to some weighting scheme, possibly different from α_i , $i = 1, 2, \dots, N$. It is clear that this method generalizes the risk-neutral case. However, we do not follow this method because it does not guarantee Pareto-optimality of the final outcome.

7. Conclusion and Further Research

We have proposed a definition of coordination of a supply chain consisting of risk-averse agents. We show that, to coordinate such a chain, the first step is to characterize the set of Pareto-optimal solutions and select a solution from this set based on the agents' respective bargaining powers. The second step is to design a contract to achieve the selected solution.

In the risk-neutral case, it is easy to see that an action pair is Pareto-optimal if and only if the supply chain's expected profit is maximized. But in the risk-averse case, it is more difficult to find Pareto-optimal actions. We characterize Pareto-optimal solutions in

Figure 1 Example of a Pareto-optimal frontier.



three specific cases of a supply chain involving one supplier and one retailer, and in each case we design a flexible contract to coordinate the channel. Furthermore, we discuss the bargaining issue in one of the cases. We provide answers to the following questions: What is the optimal external action of the supply chain and what is the optimal sharing rule?

In the specific cases that we have considered, we are able to obtain Pareto-optimal actions by first obtaining a Pareto-optimal sharing rule that can be used with any external action. This property allows us to obtain an objective function for the supply chain, whose optimization yields an external action, which together with the sharing rule provides us with Pareto-optimal solutions.

In more general cases, however, we do not have the above property, and therefore, the sequential procedure used in the special cases does not work. In such cases, we show by a specially constructed example, that obtaining Pareto-optimal solutions requires finding first the Pareto-optimal sets corresponding to external actions, and then identifying the Pareto-optimal frontier of the union of these sets. Moreover, the standard contract forms that work for risk-neutral cases do no longer coordinate, and research is required to find new coordinating contract forms.

Acknowledgments

We are grateful to Yehuda Bassok, Metin Cakanyildirim, Qi Feng, Ruihua Liu, Harry M. Markowitz, Candace Yano, Hong Yin, Hanqin Zhang, and Jing Zhou for helpful discussions, comments and suggestions. We also thank the AE and the referees for their constructive suggestions.

Editor's note: This paper was one of the two papers that won the 2004 Wickham Skinner Best Paper Award.

References

- Agrawal, V., S. Seshadri. 2000a. Impact of uncertainty and risk aversion on price and order quantity in the newsvendor problem. *Manufacturing and Service Operations Management* 2(4) 410–423.
- Agrawal, V., S. Seshadri. 2000b. Risk intermediation in supply chains. *IIE Transactions* 32 819–831.
- Arrow, K. J. 1951. *Social choice and individual values*. Wiley, New York, New York.
- Bouakiz, M., M. J. Sobel. 1992. Inventory control with an expected utility criterion. *Operations Research* 40(3) 603–608.
- Buzacott, J., H. Yan, H. Zhang. 2001. Optimality criteria and risk analysis in inventory models with demand forecast updating, Working paper, The Chinese University of Hong Kong, Shatin, Hong Kong.
- Cachon, G., M. Lariviere. 2002. Supply chain coordination with revenue-sharing contracts: strengths and limitations, Working paper, University of Pennsylvania, Philadelphia, Pennsylvania.
- Cachon, G. 2002. Supply chain coordination with contracts, Working paper, University of Pennsylvania, Philadelphia, Pennsylvania.
- Chen, F., A. Federgruen. 2000. Mean-variance analysis of basic inventory models, Working paper, Columbia University, New York, New York.
- Chen, X., M. Sim, D. Simchi-levi, P. Sun. 2003. Risk aversion in inventory management, Working paper, MIT, Cambridge, Massachusetts.
- Chopra, S., P. Meindl. 2001. *Supply chain management*. Prentice-Hall, Englewood Cliffs, New Jersey.
- Eeckhoudt, L., C. Gollier, H. Schlesinger. 1995. The risk averse (and prudent) newsboy. *Management Science* 41(5) 786–794.
- Eliashberg, J., R. L. Winkler. 1981. Risk sharing and group decision making. *Management Science* 27(11) 1221–1535.
- Gan, X., S. Sethi, H. Yan. 2003. Channel coordination with a risk-neutral supplier and a downside-risk-averse retailer, Working paper, University of Texas at Dallas, Richardson, Texas.
- Gaur, V., S. Seshadri. 2003. Hedging inventory risk through market instruments, Working paper, New York University, New York, New York.
- Harsanyi, J. C. 1955. Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of Political Economy* 63 309–321.
- Harsanyi, J. C. 1977. *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*. Cambridge University Press, Cambridge.
- Jorion, P. 2001. *Value at Risk*. McGraw-Hill, New York, NY.
- Lavalle, I. H. 1978. *Fundamentals of Decision Analysis*, Holt, Rinehart and Winston, New York, NY.
- Lau, H. S. 1980. The newsboy problem under alternative optimization objectives. *Journal of Operations Research Society* 31(5) 393–403.
- Lau, H., A. H. L. Lau. 1999. Manufacturer's pricing strategy and return policy for a single-period commodity. *European Journal of Operations Research* 116 291–304.
- Levi, H., H. Markowitz. 1979. Approximating expected utility by a function of mean and variance. *American Economic Review* 69(3) 308–317.
- Markowitz, H. 1959. *Portfolio selection: Efficient diversification of investment*. Cowles Foundation Monograph 16, Yale University Press, New Haven, Connecticut.
- Nagarajan, M., Y. Bassok. 2003. A bargaining framework in supply chains. Working paper, University of Southern California, Los Angeles, California.
- Nash, J. F. 1950. The bargaining problem. *Econometrica* 18 155–162.
- Pasternack, B. 1985. Optimal pricing and returns policies for perishable commodities. *Marketing Science* 4(2) 166–176.
- Pratt, W. J. 1964. Risk aversion in the small and in the large. *Econometrica* 32 122–136.
- Raiffa, H. 1970. *Decision analysis*. Addison-Wesley, Reading, Massachusetts.
- Schweitzer, M. E., G. P. Cachon. 2000. Decision bias in the newsvendor problem with a known demand distribution: Experimental evidence. *Management Science* 46(3) 404–420.
- Szegö, G. (Eds.). 2004. *New risk measures for the 21st century*. Wiley, West Sussex, England.
- Tayur, S., R. Ganeshan, M. Magazine (Eds.). 1999. *Quantitative models for supply chain management*. Kluwer, Boston, Massachusetts.
- Tsay, A. 2002. Risk sensitivity in distribution channel partnerships: Implications for manufacturer return policies. *Journal of Retailing* 78(2) 147–160.
- Van Mieghem, J. A. 2003. Capacity management, investment, and hedging: Review and recent developments. *Manufacturing and Service Operations Management* 5(4) 269–302.
- Van Neumann, J., O. Morgenstern. 1944. *Theory of games and economic behavior*. Princeton University Press, Princeton, New Jersey.
- Wilson, R. 1968. The theory of syndicates. *Econometrica* 18 155–162.