Technical Memorandum Number 823

Risk Aversion and Supply Chain Contract Negotiation

by

Matthew J. Sobel
Danko Turcic

December 2008

Department of Operations
Weatherhead School of Management
Case Western Reserve University
330 Peter B. Lewis Building
Cleveland, Ohio 44106

Risk Aversion and Supply Chain Contract Negotiation

Matthew J. Sobel • Danko Turcic

December 3, 2008

Department of Operations, Weatherhead School of Management, Case Western Reserve University, Cleveland, OH 44106

Olin Business School, Washington University in St. Louis, St. Louis, MO 63130

matthew.sobel@case.edu • *turcic@wustl.edu*

We employ the Nash bargaining solution to determine the value of an opportunity to negotiate a contract in an archetypal supply chain game. This unifies the allocation of payoffs and the selection of the type of contract in a bilateral supply chain; a supplier delivers goods to a newsvendor retailer who stocks in anticipation of stochastic seasonal demand. Both firms have concave, ordinal utility functions which allows us to model both risk-neutrality and risk-aversion. We characterize the Nash-optimal contracts, and contrast the cases where both firms are risk-neutral and where at least one of the firms is risk-averse. Risk aversion, it turns out, has a significant impact on the contract terms and one firm's risk aversion may be advantageous to the other. This suggests that each firm should strive to find a supply chain partner based on its own and its prospective partner's sensitivities to risk. For example, risk-neutral firms have an incentive to *avoid* other risk-neutral firms as their supply chain partners.

---

1. Introduction

Negotiation is meaningful only if neither party can enforce a take-it-or-leave-it decision on the other; consequently, the value of a negotiation opportunity is ambiguous. The value of a supply chain contract negotiation is ambiguous too because the end-user demand is uncertain, the contract's structure and parameters will be the crux of the negotiation, and the other negotiating firms' decision rules are unknown. The first consideration is a market risk, and the second and third are strategic risks. The supply chain research literature often evaluates market risk with risk-neutrality; so it replaces random monetary payoffs with expected values. The game theory approach to supply chain strategic risk has maximized the overall "size of the pie." It has been common to assume that there is a given type of contract and one of the firms acts as the Stackelberg "leader" while the other has no recourse but to accept its leadership. Nagarajan and Bassok (2008), however, observe that " . . . Stackelberg games endow players with disproportionate negotiation power . . . there are environments where negotiation power is not as discrete or extreme as dictated by a Stackelberg game . . . ." When both firms have equal negotiating power, little is known about (a) the selection of type of contract, (b) dividing the pie among the firms, (c) the effect of risk-averse preferences, and (d) the negotiation of endogenous prices. This paper provides a unified approach to these issues.

There are at least two reasons for the lack of research attention to dividing the pie. First, the leader in the standard Stackelberg game approach has the power to compel the other firm to accept its leadership, and the follower is resigned to reacting to the leader's decision as advantageously as possible. Of course, the leader's decision extracts all of the supply chain's profit above the follower's reservation level. That is, the Stackelberg game solution yields a unique division of the pie. Another reason for the lack of attention is that it seems more important to enlarge the pie than to dwell on the division of a smaller pie; a larger pie permits every member firm to have a bigger piece. As a result, the literature is more advanced in its analysis of behavior than of value. When neither firm has the power to enforce a leader's role, we note in §2 that the game yields a set of equilibria with a wide range of payoffs. This status is unsatisfactory because the perceived value of an opportunity to negotiate a supply chain contract would affect a firm's decision on whether or not to proceed and, if so, its subsequent decisions.

Shubik (1982) (p.179) observes that "On many occasions ... it is important to have a specific "one-point" solution - a single payoff vector that expresses the value of the game to each of the players. This is desirable not only for the intellectual satisfaction of coming up with a clear, sharp answer to a complex problem. It is desirable also because many applications of game theory ... by their very nature seem to demand a straight answer ... to the question of what a particular competitive position is actually worth."

It is well-known that there are many efficient Pareto-improving contracts, i.e., many contracts maximize the overall size of the pie in the bilateral supply chain coordination model studied by Pasternack (1985) (and subsequently by many others). Throughout the paper we refer to results as "well-known" if they are included in expositions by Cachon (2003) or Cachon and Netessine (2004). Well-known results treat profit uncertainty with risk-neutrality and utilize the *equilibrium point* (or *Nash equilibrium*) solution concept.

Risk posture

If the retailer and the supplier are risk-neutral, then maximizing the size of the pie corresponds to maximizing the aggregate supply chain profit. If at least one of the players maximizes its expected utility of profit, and its utility function is not entirely linear, then the well-known efficiency properties in the risk-neutral case are no longer valid. Nevertheless, the Stackelberg game is well-defined. The leader accurately anticipates that the follower would react to the leader's decision by maximizing the expected utility of its own profit. So the leader factors the effect of the follower's self-interested response on its own expected utility, and makes a choice that maximizes the expected utility of its own profit. How efficient is this solution?

If at least one of the players is risk-averse, the "size of the pie" is ill-defined. The sum of the expected utilities of the retailer's and supplier's net profits does *not* correspond to the size of the pie for the following reason. Utility functions convey only *ordinal* information – a positive affine transformation of a utility function is again a utility function. That is, if $u(\cdot)$ is a utility function, so also is $au(\cdot) + b$ for all $a > 0$ and b. As a consequence, if $u_r(\cdot)$ and $u_s(\cdot)$ are the retailer's and supplier's utility functions, then the set of sums of positive affine transformations of their utilities includes $\alpha u_r(\cdot) + (1 - \alpha)u_s(\cdot)$ for all $0 < \alpha < 1$. So the set of sums of their utilities includes criteria that give overwhelming weight to the preferences of the supplier, and other criteria that give overwhelming weight to the preferences of the retailer.

The importance of considering risk-aversion in a supply chain may not be obvious. After all, there are well-known diversification arguments that support the commonly made risk-neutrality assumption. Van Mieghem (2003) notes that the question of whether firms should care about risk amounts to determining the objective of the firm. For privately held firms, the objective of an owner whose assets are tied up in a firm is to maximize the owner's expected utility of profit. For publicly held firms operating in an imperfect market, Van Mieghem (2003) notes that "non-linearities enter into the objective and firms do care about risk." In addition to market imperfections, there are principal agent phenomena that cause the risk-neutrality assumption to break down; we refer the reader to Pindyck and Rubinfeld (2008) for further discussion. Empirical studies find that firms are not effective at hedging supply chain risk (c.f. Williams (2008)). Therefore, it is important to develop supply chain contracting results without assuming risk-neutrality.

Pure bargaining problem

In game theory it is not uncommon to find several different solution concepts that build on the same descriptive formulation (see Shubik (1982) for further discussion). The model in this paper preserves the economics of the supplier/retailer interaction underlying the well-known results. However, it uses the *Nash bargaining solution* to determine the value of a competitive position in the archetypal supply chain game. In that game, a wholesaler supplies goods to a newsvendor retailer who stocks in anticipation of stochastic seasonal demand. If the players agree on the type of contract and its terms, then the contract specifies their payments (usually subject to market risk). If the players fail to reach agreement, then their payoffs are disagreement amounts. This is a structured two-person pure bargaining problem.

Generally, the *pure bargaining problem* is the paradigm "in which a group of two or more participants is faced with a set of feasible outcomes, any one of which will be the result if it is specified by the unanimous agreement of all the participants. In the event that no unanimous agreement is reached, a given disagreement outcome is the result [Roth (1979)]." Cachon and Netessine (2004) and Nagarajan and Sošić (2008) emphasize the noncooperative and cooperative branches of game theory, respectively, as they analyze, discuss, and survey supply chain games. Bargaining as a branch of game theory begins with the seminal contributions by Nash (1950, 1953).

A *strategy* in the pure bargaining problem specifies the decision rule that each player will use. As we explain in §2, the *Nash solution* of the pure bargaining problem is obtained as follows. For each strategy and each player, consider the difference between that player's utility if that strategy is employed and that player's utility of the disagreement outcome. The Nash solution is a strategy that maximizes the product of the differences across all players. So the Nash solution depends neither on risk-neutrality nor on maximizing the sum of the bargainers' utilities. This "solution" is employed widely in economics, and it has been advocated at least since Zeuthen (1930). We too employ the Nash bargaining solution.

Although there are non-Nash solutions of the pure bargaining problem, the Nash solution is employed in most instances. See Myerson (2004), Osborne and Rubinstein (1994), Roth (1979), and Shubik (1982) for an array of game theoretic approaches to bargaining. See Nagarajan and Sošić (2008) for justifications of the axiomatic approach in Nash (1950) in the analysis of supply chain contracts. This paper is closely related to Nagarajan and Sošić (2008) who evaluate a special case of our model, namely a buy-back contract between a risk-neutral supplier and a risk-averse retailer.

Some of the earlier applications of bargaining models in supply chain, inventory, and capacity management include Kohli and Park (1989) who study the implementation of quantity discounts in an inventory problem, Reyniers and Tapiero (1995) who analyze a quality control problem, and Chod and Rudi (2006) who study investments in capacity. The value of the option to subcontract is studied in Van Mieghem (1999). To our knowledge, this is the first use of a bargaining framework in supply chain coordination and contract selection. Gurnani and Shi (2006) study outcomes in a bilateral supply chain where the supplier is unable to guarantee a delivery. In this setting they derive a price of a supply contract and the Nash-optimal order quantity. The uncertainty in the model is the supplier's unreliability and most results are obtained under risk-neutrality. They also show a risk-averse solution using polynomial utility functions. Dong and Liu (2005) study a special type of supply chain, forward supply contracts on non-storable commodities, where the supplier and the retailer are risk-averse with mean-variance objectives. Wu (2004) exploits the Nash bargaining solution to study the formation of intermediaries in supply chains. Most recently, Nagarajan and Bassok (2008) use a bargaining framework to study a two-stage supply chain with one risk-neutral buyer sequentially negotiating with a series of potential risk-neutral vendors who can form coalitions.

Contributions

This seems to be the first analysis of the bilateral supply chain model with risk-averse players having general ordinal utility functions. In contrast, Gan et al. (2004) exemplifies the insertion of *cardinal* utility functions in the equilibrium point approach. We use the Nash bargaining solution concept to provide a unified approach to the problem of allocating payoffs and selecting the type of contract in a bilateral supply chain with risk-averse players. The results enable the upstream and the downstream firms to value an opportunity to negotiate a supply chain contract. We show that risk-aversion can have a significant impact on the terms of the contract and that both firms have strong incentives to choose a supply chain partner based on its sensitivity to risk. For example, we identify conditions under which the expected profit of a risk-neutral upstream firm increases with risk-aversion of the downstream firm. Therefore, risk-neutral firms may be motivated to *avoid* other risk-neutral firms as their supply chain partners. Similarly, our results also reveal that risk-averse firms are better off partnering with either risk-neutral firms or firms that are extremely risk averse.

In §3 we show that the Nash bargaining approach is powerful even if preferences are risk-neutral. This complements Nagarajan and Sošić (2008) with a general supply chain bargaining game of complete information in which the supplier and the retailer negotiate (1) contract type; (2) contract parameters; and (3) order quantity. In §4 we compare the risk-neutrality results with the solution when the players are risk-averse. It is well-known that a degraded supply chain performance (when compared to the risk-neutral theoretical outcome) is commonly observed in practice. However, it has been conjectured that this is not necessarily due to incompetent managers or naive operating policies. Using Nash's bargaining model, we show how a dominant supplier (or retailer) may have an incentive to reduce the performance of the integrated channel in exchange for its own gain! A numerical example in §5 illustrates the results in §3 and §4. The retail price is exogeneous in §3 and

§4, but it is endogenous in §6 which generalizes the model by permitting the probability distribution of retail demand to depend on a price that may be influenced by the supplier as well as the supplier. The next section specifies the model and formally describes the game between the supplier and the retailer.

2. The Model

Consider an archetypal supply chain with an upstream supplier and a newsvendor retailer who stocks a single product in anticipation of stochastic seasonal demand, D. The supplier produces q items after receiving the retailer's order for q items and delivers q items at the beginning of the selling season. There is no risk of failure to deliver due to credit failure or any other reason.

We use s and r to label the supplier and the retailer, also referred to as *players*, and they are assumed to have preferences on the set of payoffs of this supply chain that are consistent with the existence of utility functions, say $u_s(\cdot)$ and $u_r(\cdot)$. We assume that both utility functions are concave, strictly increasing, and twice differentiable. After §3 we do *not* assume that the utility functions are linear; linearity corresponds to risk-neutrality. Both players are assumed to be self-serving expected utility maximizers.

We adopt the following notation: F is the distribution function of demand; p is the exogenous retail price; c_s and c_r are the supplier's and the retailer's unit production costs; g_s and g_r are the supplier's and the retailer's unit shortage penalty costs; and v is the unit net salvage value of unsold goods. Let $c = c_r + c_s$ and $g = g_r + g_s$; we assume that $v < c < p$.

In order to do business with each other, the supplier and the retailer need to agree formally on a number of details that are typically specified in a supply contract. Here, these details are (1) a type of contract; (2) contract parameters; and (3) an order quantity, q.

Thanks to the large literature on supply contracts, diverse contract arrangements are available to the players and our analysis does *not* a priori impose a particular type. Here, *the players choose* a type of contract and the parameters of that type. Some of the well-known types are wholesale, buy-back, revenue-sharing, and quantity discount contracts (cf. Pasternack (1985), Cachon and Lariviere (2005), Tomlin (2003), and Cachon (2003)). For example, if the players agreed to use a buy-back contract, then they would also select the wholesale and buy-back prices as contract parameters.

Since both players are expected utility maximizers and the model assumes voluntary compliance, both players need to agree on a supply order quantity, q. That is, the self-serving retailer must be willing to order q units and the self-serving supplier must be willing to supply them. We study the contract types, contract parameters, and order quantities that the utility maximizing supplier and retailer choose in this setting.

The model is transformed into a game with the following *descriptive terminology*: (i) the supplier and the retailer are *players*; (ii) the combinations of order quantity, q, supply contract type and contract parameters are *strategies*; and (iii) the supplier's and the retailer's expected utilities of profit are *payoffs*.

The retailer's and the supplier's profit are the following functions of the stochastic demand, D:

$$\pi_r = pS + vI - c_r q - g_r L - \sum_{k \in K} i_k T_k = (p - v + g_r)S - (c_r - v)q - g_r D - \sum_{k \in K} i_k T_k, \quad \text{(1a)}$$

$$\pi_s = g_s S - c_s q - g_s D + \sum_{k \in K} i_k T_k, \quad \text{(1b)}$$

where, respectively, $S = \min\{q, D\}$, $I = (q - D)^+$, and $L = (D - q)^+$ denote the sales quantity, excess inventory, and excess demand (lost sales). Also, T_k is a transfer payment function specific to a given contract type, i_k is an indicator variable denoting that a particular type of supply contract is in force, and K is a discrete set of contract types.

To illustrate the role of the set K, suppose the players choose among buy-back, revenue-sharing, and quantity discount contracts. Label buy-back with $k = 1$, revenue-sharing with $k = 2$, and quantity-discount with $k = 3$. So $K = \{1, 2, 3\}$. If the players agree on a revenue-sharing contract, say, to write the supplier's and the retailer's profits in (1a) and (1b), let $i_2 = 1$ and $i_1 = i_3 = 0$. The corresponding transfer payments, T_k, are given by:

$$\text{Buy-back:} \qquad T_1 = bS + (w_b - b)q, \quad \text{(2a)}$$

$$\text{Revenue-sharing:} \qquad T_2 = [w_r + (1 - \phi)v]q + (1 - \phi)(p - v)S, \quad \text{(2b)}$$

$$\text{Quantity discount:} \qquad T_3 = w_q q, \quad \text{(2c)}$$

The contract parameters in this example are the wholesale price, w, the buy-back price, b, and the fraction of sales revenue retained by the retailer, ϕ. Thus, the retailer's and the supplier's profits under revenue-sharing can be written as

$$\pi_r = pS + vI - c_r q - g_r L - i_2 T_2 = (\phi(p - v) + g_r)S - (w_r + c_r - \phi v)q - g_r D,$$

$$\pi_s = g_s S - c_s q - g_s D + i_2 T_2 = [w_r + (1 - \phi)v - c_s]q + [(1 - \phi)(p - v) + g_s]S - g_s D,$$

Diverse game theoretic and decision theoretic solution concepts can be applied to the preceding description. For example, we define the worth of a contract to player j ($j \in \{r, s\}$) as j's certainty equivalent of the payoff under the contract.

Definition 1. p_j is the *worth of a supply contract* to player j if $u_j(p_j) = E[u_j(\pi_j)]$, $j \in \{r, s\}$.

In words, p_j is the worth of a contract to player j if j finds the contract's random payoff π_j and the certain amount p_j equally attractive.

§1 alludes to the ambiguity of the profit allocations between the retailer and the supplier in the following example. Let $p = \$25, c_r = \$5, c_s = \$10, v = \$5$, $g_r = g_s = \$0$, and let demand, D, be uniformly distributed on $[0, 100]$. Assume both players have a reservation expected profit level of \$2.50. The Stackelberg game with a buy-back contract yields an optimal order quantity $q^* = 100$ units and a corresponding expected total supply chain profit of \$250. If the supplier is the Stackelberg leader then the expected profit allocations to the retailer and the supplier, respectively are the components of $\{\$2.50, \$247.50\}$. If neither player is the leader and their choices are made simultaneously, then many pairs of wholesale and buy-back prices, $\{w, b\}$, achieve the expected total supply chain profit of \$250 including $\{\$19.90, \$19.80\}$, $\{\$17.50, \$15.00\}$, $\{\$15.00, \$10.00\}$, $\{\$12.50, \$5.00\}$, and $\{\$10.1, \$0.20\}$. However, these pairs result in the following expected profit allocations to the retailer and the supplier, respectively: $\{\$2.50, \$247.50\}$, $\{\$62.50, \$187.50\}$, $\{\$125.00, \$125.00\}$, $\{\$187.50, \$62.50\}$, $\{\$247.50, \$2.50\}$. All of these pairs of expected profit allocations are consistent with equilibrium point payoffs (see §3 in Rubinstein (1982) for further discussion).

In the supply chain bargaining game, let M be a non-empty compact set of pure (i.e., non-randomized) strategies, and let m denote an element in M. So m is a vector that consists of the order quantity, q; a vector of zero-one variables that indicate the type of contract, $i = (i_1, i_2, \dots)$, where $i_1 + i_2 + \cdots = 1$; and a vector of contract parameters, $t = (w_b, b, w_r, \phi, w_q, \dots)$.

Neither player, however, is restricted to pure strategies and may employ a mixed (randomized) strategy. Let $\Delta(M)$ denote the set of probability distributions over the set of Borel measurable subsets of M. The players' pure strategy monetary profits, given by (1), are random variables due to the randomness of demand, and are denoted $\pi_s(m)$ and $\pi_r(m)$. Let π_j^γ denote j's random profit if the players employ the (possibly) randomized strategy γ. Therefore, π_j^γ is the following mixture of pure strategy profits:

$$\pi_j^\gamma = \int_M \pi_j(m) d\gamma(m), \ \gamma \in \Delta(M), \ j \in \{r, s\}. \tag{3}$$

Let U' be the set of all pairs of expected utilities that the players can achieve with pure strategies, and let U denote the set of such pairs when the players resort to randomized strategies:

$$U' = \{(Eu_s[\pi_s(m)], Eu_r[\pi_r(m)]) : m \in M\} \quad \text{and} \quad U = \{(Eu_s(\pi_s^\gamma), Eu_r(\pi_r^\gamma)) : \gamma \in \Delta(M)\}.$$

Both sets can be displayed graphically in the Cartesian plane where it is apparent that U is convex and compact. Convexity is intuitive because a point lying on a straight line that connects two points in U can be achieved by an appropriate choice of a mixed strategy. Therefore, $U' \subseteq U$. Lemma 4 in §B shows that $U \subseteq U'$ for i fixed (meaning for a given type of supply contract), so $U = U'$. Thus, for a given contract type, the players need not randomize to choose the contract parameters. The concavity of $\pi_j(\cdot), j \in \{r, s\}$ on M (for each outcome in the sample space of D) is a sufficient condition for this property.

The following sequence of moves occurs in this game. First, the supplier and the retailer choose mixed strategies called *threats*, say τ_s and τ_r, that they will use if they cannot agree on a supply contract. The threats may be invoked during the contract negotiation process and have associated monetary profits that are labeled $\pi_s^{\tau_s}$ and $\pi_r^{\tau_r}$. These payoffs need not be deterministic. For example, a threat strategy may consist of walking away when faced with an expected utility level that is lower than is available via a competing supply contract. We assume that both players know their own and each other's utility functions, threat strategies.

Next, the supplier and the retailer *independently* select expected utility levels, labelled u_s^δ and u_r^δ, which they demand from a potential supply chain contract. Since a positive affine transformation of a utility function is again a utility function, the particular scalings of $u_s(\cdot)$ and $u_r(\cdot)$ precede the selection of u_s^δ and u_r^δ. Operationally, we associate u_s^δ and u_r^δ with (possibly) mixed strategies δ_s and δ_r that achieve the expected utility levels, i.e., $Eu_j(\pi_j^{\delta_j}) = u_j^\delta$, $j \in \{r, s\}$.

The game ends with a determination of payoffs. If there is a mixed strategy $\gamma \in \Delta(M)$ such that $Eu_s(\pi_s^\gamma) \geq u_s^\delta$ and $Eu_r(\pi_r^\gamma) \geq u_r^\delta$, then the supplier and the retailer receive their *demanded* expected utility levels (u_s^δ, u_r^δ). Operationally, we construe this to mean that they receive the profits $(\pi_r^{\delta_r}, \pi_s^{\delta_s})$. Otherwise, threats are carried out and the players attain $(Eu_s(\pi_s^{\tau_s}), Eu_r(\pi_r^{\tau_r}))$. Clearly, given this payoff function, both players have the incentive to demand as much expected utility as possible without losing compatibility.

Two of the justifications for what is now called the *Nash solution* of the bargaining problem are due to Nash who shows that a necessary and sufficient condition for a (possibly) randomized policy $\gamma \in \Delta(M)$ to be a solution is that $Eu_j(\pi_j^\gamma) \geq Eu_j(\pi_j^{\tau_j})$, $j \in \{r, s\}$, and γ achieves

$$\max_{\gamma \in \Delta(M)} \left[Eu_s(\pi_s^\gamma) - Eu_s(\pi_s^{\tau_s})\right]\left[Eu_r(\pi_r^\gamma) - Eu_r(\pi_r^{\tau_r})\right]. \tag{4}$$

Note that the solution set of (4) is invariant with respect to positive affine transformations of $u_s(\cdot)$ and $u_r(\cdot)$. Henceforth, we consider strategies only if they are at least as good as threats, i.e., $Eu_j(\pi_j^\gamma) \geq Eu_j(\pi_j^{\tau_j})$, $j \in \{r, s\}$. So we refer to (4) alone as a necessary and sufficient condition.

3. Risk-Neutrality

In this section, we obtain the following results under the assumption that both players are risk-neutral. The Nash bargaining solution assigns a unique expected profit to each player which increases with the expected payoff from the player's threat strategy. If the retailer and supplier have the same expected payoffs from their threat strategies, then the Nash solution assigns them equal expected profit. If the players have concave expected profit functions, then the Nash-bargaining-optimal order quantity in (4) is non-randomized, independent of the supply contract type, and it maximizes the expected profit of the entire supply chain. Concavity also implies that both players randomize, if at all, only to choose the type of supply chain contract; the choices of order quantity and contract parameter are non-randomized.

We assume that the utility functions $u_s(\cdot)$ and $u_r(\cdot)$ are increasing and linear and that the retailer's and supplier's profit functions, $\pi_r(\cdot)$ and $\pi_s(\cdot)$, are concave on M (for each outcome of demand). The linearity of $u_s(\cdot)$ and $u_r(\cdot)$ permits the necessary and sufficient condition for a Nash bargaining solution, namely (4), to be written as follows:

$$\max_{\gamma \in \Delta(M)} \left[E\pi_s^{\gamma} - E\pi_s^{\tau_s}\right]\left[E\pi_r^{\gamma} - E\pi_r^{\tau_r}\right].$$

Here, the players may randomize their choice of order quantity, q, type of supply contract, and contract parameters.

There is a further simplification of (4) due to the concavity of $\pi_r(\cdot)$ and $\pi_s(\cdot)$ on M (for each outcome of demand) and to the following result that is proved in §A.

Lemma 1. *If the utility functions $u_j(\cdot)$ are concave increasing and the profit functions $\pi_j(\cdot)$, $j \in \{r, s\}$, are concave on M (for each outcome of demand), then there is no loss of optimality in (4) if the supplier and the retailer randomize at most only in their choice of the type of supply chain contract.*

So we assume that the supplier and the retailer make deterministic choices of the order quantity and the contract parameters, and they randomize (at most) in their choice of the type of contract. This permits the following further simplification of (4):

$$\max \left[E(E\pi_s(m)|i) - E\pi_s^{\tau_s}\right]\left[E(E\pi_r(m)|i) - E\pi_r^{\tau_r}\right]. \tag{5}$$

The maximization here is with respect to $m \in M$ and the probabilities of choosing various types of contracts. This leads to a further simplification of the Nash bargaining solution.

Proposition 1. *Without loss of optimality in (5), the players choose supply contract parameters that achieve equal gains in their expected profits over the expected payoffs of their threat* *strategies and randomize arbitrarily among the types of supply contracts that can achieve that payoff allocation.*

This result has an equivalent statement in terms of arithmetic means rather than equal gains.

Corollary 1. *Without loss of optimality in (5), the players choose supply contract parameters that maximize the arithmetic mean of the expected profit gains over the expected payoffs from their threat strategies, and randomize arbitrarily among contract types parametrized to achieve such a payoff allocation.*

The next result, the advantage of having a threat strategy with a high expected monetary profit, follows immediately from Proposition 1.

Corollary 2. *Under risk-neutrality, for either player, a higher expected threat strategy payoff than the opponent's translates into capturing a greater portion of the total supply chain profit.*

These two results indicate that the Nash solution specifies a unique allocation of payoffs to the retailer and the supplier, and highlight how each player can improve its bargaining position and increase its expected profit from a particular supply chain deal. In particular, Corollary 2 implies that when both players are risk-neutral and employ threat strategies with equal expected payoffs, both players should expect to "walk away" with equal expected profit.

Suppose that x contract types can achieve the equal gains over threats allocations in (5). (It is well-known that some contracts are unable to achieve all payoff allocations; for example, a quantity discount contract has an upper bound on the allocation of expected profit to the retailer. Without loss of optimality, the players avoid such contracts in (5). See Cachon (2003) for a survey of contracts able to achieve arbitrary payoff allocations.) As Proposition 1 indicates, the players may arbitrarily choose one of these x types or they may randomize their choice. Some professional practice is consistent with this finding; *The Wall Street Journal* (2006) reports that Wal-Mart maintains many different contract arrangements and emphasizes open and transparent relationships with its suppliers "to have the ability to work things out as situations are not always clear cut." Using the definitions of $\pi_r(m)$ and $\pi_s(m)$ in (1), the total supply chain profit is

$$\pi(q) = \pi_r(m) + \pi_s(m) = (p - v + g)S - (c - v)q - gD. \tag{6}$$

Proposition 2. *The non-randomized order quantity that is optimal in (5) maximizes the expected value of the total supply chain profit, $E[\pi(q)]$, and is given by*

$$q^* = F^{-1}\left(\frac{p-c+g}{p-v+g}\right). \tag{7}$$

This proposition reveals that the players should choose a supply quantity q^*, regardless of the type of supply contract, that maximizes the expected profit of the entire supply chain. Expression (7) is the well-known equilibrium order quantity. Therefore, the equilibrium point approach in the Stackelberg game and the Nash bargaining solution are complementary and each clarifies the other. The following proposition summarizes these points.

Proposition 3. *In the Nash bargaining solution, the players choose a non-randomized order quantity, given by (7), that maximizes the expected total profit of the supply chain. The players choose a type of supply contract by randomizing among contracts that are parametrized to achieve equality among players of the gain in expected profit over the expected monetary payoff from the threat strategy.*

Since the Nash bargaining solution assigns a unique expected monetary payoff to each player, the supplier and the retailer can evaluate the value of an opportunity to enter into a supply chain contract; the next result follows from Definition 1 of the worth of a supply contract.

Remark 1. *If both players are risk-neutral, then the worth of a potential supply chain contract to player $j \in \{r, s\}$ in the Nash bargaining solution is the maximal expected profit achieved in (5) with the potential supply contract i.*

The Nash solution (5) is further simplified if the set of feasible contracts is constrained. We illustrate the simplification with a reduction to buy-back and revenue-sharing contracts with transfer payment functions given by (2a) and (2b). Since, under risk-neutrality, these contracts are equivalent, it is well-known that there is a value of λ, $0 \le \lambda \le 1$, such that $\pi_s(m)$ and $\pi_r(m)$ given by (1) can be written as

$$\pi_r(m) = \lambda\pi(q) + D\lambda g - Dg_r, \tag{8a}$$

$$\pi_s(m) = (1-\lambda)\pi(q) - D\lambda g + Dg_r, \tag{8b}$$

where $\pi(q)$ is given by (6). So λ is a parameter that allocates the total supply chain profit between the supplier and the retailer; the retailer receives fraction λ of the total supply chain profit and the supplier receives the remaining $1-\lambda$.

In this case, the Nash bargaining solution (5) is

$$\max_{\substack{\lambda\in[0,1]\\ q\in m}} \left[(1-\lambda)E\pi(q) - \mu\lambda g + \mu g_r - E\pi_s^{\tau_s}\right]\left[\lambda E\pi(q) + \mu\lambda g - \mu g_r - E\pi_r^{\tau_r}\right], \tag{9}$$

where $\mu = E(D)$, the expected value of demand. In this maximization, the first-order conditions with respect to q and λ are

$$\lambda\frac{\partial E\pi(q)}{\partial q} = (1-\lambda)\frac{\partial E\pi(q)}{\partial q} = 0 \qquad \Leftrightarrow \qquad \frac{\partial E\pi(q)}{\partial q} = 0 \tag{10}$$

and

$$\left[E\pi(q) + \mu g\right]\left[2\lambda E\pi(q) - E\pi(q) + E\pi_s^{\tau_s} - E\pi_r^{\tau_r} - 2\mu g_r + 2\lambda\mu g\right] = 0. \tag{11}$$

An order quantity $q = q^*$, given by (7), solves (10) with $\lambda = \lambda^*$ where

$$\lambda^* = \frac{E\pi(q) + E\pi_r^{\tau_r} - E\pi_s^{\tau_s} + 2\mu g_r}{2[E\pi(q) + \mu g]}, \tag{12}$$

solves (11). This confirms the structural results obtained earlier (cf. Proposition 3); the players choose an order quantity that maximizes the expected profit of the entire supply chain and, by selecting $\lambda = \lambda^*$, they partition the total expected profit so as to ensure equal gains over the expected threat payoffs, $E\pi_s^{\tau_s}$ and $E\pi_r^{\tau_r}$. For the special case when $g_r = g_s = 0$ and $E\pi_r^{\tau_r} = E\pi_s^{\tau_s}$, implying that both players are utilizing equally effective threat strategies, (11) yields $\lambda^* = 1/2$, implying that the total supply chain profit would be split exactly in half.

4. Risk-Aversion

In this section we assume that the utility functions $u_s(\cdot)$ and $u_r(\cdot)$ are concave, increasing, and twice differentiable, at least one of them is non-linear, and the retailer's and the supplier's profit functions, $\pi_r(\cdot)$ and $\pi_s(\cdot)$, are concave on M for each outcome of demand. We show that risk-aversion has the following consequences. The order quantity in the Nash bargaining solution is lower than under risk-neutrality, and this causes a decrease in the expected total supply chain profit. The division of the total profit between the supplier and the retailer is asymmetric if they have different sensitivities to risk. Risk-aversion can be either an advantage or a disadvantage to a player's payoff. We provide sufficient conditions for a player's expected utility of profit to increase (or decrease) as the other player's risk-aversion increases. One implication of this result is that the supplier or the retailer can increase its expected utility of profit even as the expected total profit of the supply chain decreases. Another implication is that a player can increase its expected utility of profit in a supply chain by selecting the other player partly on the basis of its sensitivity to risk.

This section proceeds as follows. §4.1 generalizes the risk-neutral results in §3 to a setting where at least one of the players is risk averse. §4.2 explores an observation in Roth and Rothblum (1982) that risk aversion may or may not be an advantage in the bargaining game described in §2. There is a convenient transformation in §4.3 that reduces the dimensionality in the Nash bargaining solution, sharpens the general results obtained in §4.2, and leads to easily communicated sufficient conditions for one player's risk aversion to be advantageous to the other. This is followed by an illustrative example.

4.1 Preliminaries

Concave utility and profit functions simplify the solution to (4) which continues to be a necessary and sufficient condition for a Nash bargaining solution of the supply contract model. Lemma 1 simplifies (4) to the following risk-averse analog of (5):

$$\max \big\{E\big[Eu_s(\pi_s(m))|i\big] - Eu_s(\pi_s^{\tau_s})\big\}\big\{E\big[Eu_r(\pi_r(m))|i\big] - Eu_r(\pi_r^{\tau_r})\big\}. \tag{13}$$

The maximization here is with respect to $m \in M$ and the probabilities of choosing various types of supply contracts. In accordance with the lemma, for the remainder of this section we assume that the players do not randomize in their choice of the order quantity, q (which is independent of the supply contract type), or the parameters of the supply contract; they may randomize in their choice of the type of supply contract.

Since $m \in M$ and the probabilities of choosing various types of supply contracts which are optimal in (13) also maximize the square root of the maximand in (13), a mixed strategy is optimal in (13) if and only if it maximizes the geometric mean of the expected utility gains over the expected utilities from the threat payoffs. Under risk-aversion, the equal gains property (cf. Proposition 1) does not obtain. Therefore, the following Proposition is a risk-averse equivalent to Corollary 1. The proof is omitted.

Proposition 4. *Without loss of optimality in (13), the players choose supply contract parameters so as to maximize the geometric mean of the expected utility gains over the expected utilities from the threat payoffs and randomize arbitrarily among contract types parametrized to achieve such a payoff allocation.*

There is a simple formula for estimating the value of each player's bargaining position.

Remark 2. *Let $b_j(i)$ be the maximal expected utility that player j achieves in (13) with the potential type of supply chain contract i ($j \in \{r, s\}$). Then the worth of potential contract type i to player j is the certainty equivalent $\hat{p}_j(i)$ where*

$$u_j(\hat{p}_j(i)) = b_j(i),\ j \in \{r, s\}$$

The remainder of this section is presented from the supplier's point of view (for expository simplicity). However, the conclusions do not depend on player labels and they are equally valid from the retailer's point of view. Henceforth, we use the Arrow-Pratt measure of risk-aversion, namely $R_j(x) = -u_j''(x)/u_j'(x)$, $j \in \{r, s\}$, and for two real-valued functions f and g we write $f \lesssim g$ if $f(x) \leq g(x)$ for all x.

Definition 2. A player with utility function $\hat{u}(x)$ and Arrow-Pratt measure

$$\hat{R}(x) = -\hat{u}''(x)/\hat{u}'(x)$$

is *more risk-averse* than a player with utility function $u(x)$ if $R \lesssim \hat{R}$

The next set of results shows that as the retailer's risk-aversion increases, there is a decrease in the order quantity and in the expected total profit associated with the Nash bargaining solution. That is, unlike the risk-neutral setting of Proposition 2, risk-aversion diminishes the expected total profit of the supply chain.

Proposition 5. *The order quantity in the Nash bargaining solution decreases as the retailer becomes more risk-averse.*

Risk-neutral utility functions are least risk-averse among concave functions, so Proposition 5 implies that ther order quantity is smaller with risk-aversion than with risk-neutrality.

Corollary 3. *Let q'^* and q^* be order quantities in Nash bargaining solutions under risk-aversion and risk-neutrality, respectively (i.e., optimal in (5) and (13)). Then $q'^* \leq q^*$.*

Risk-aversion also diminishes the total size of the pie.

Proposition 6. *As the retailer becomes more risk-averse, the expected total supply chain profit decreases.*

4.2 Effect on the Supplier of a Change in the Retailer's Risk Aversion

The type of supply chain contract that emerges in a solution depends on the players' risk-aversions. So the results in this section transcend particular types of contracts.

In the context of a single type of contract, it seems to be well-known that risk-aversion reduces supply chain performance. This suggests that risk intermediation may always be in the best interest of both players (cf. Agrawal and Seshadri (2000)). Here we show that one player's risk-aversion may actually be advantageous in (13) to the other. That is, the supplier may be able to increase its expected utility by seeking out a *more* risk-averse retailer in spite of further reducing the expected total profit of the supply chain (cf. Proposition 6).

For expository convenience, in the remainder of this section we assume that the threat payoffs $\pi_s^{\tau_s}$ and $\pi_r^{\tau_r}$ from the threat strategies τ_s and τ_r are deterministic.

Definition 3. The retailer's monetary profit is *favorably u_r-supported at $m \in M$* if

$$u_r(\pi_r(m,d)) \geq u_r(\pi_r^{\tau_r})$$

for all outcomes d in the sample space of demand D. Otherwise it is *unfavorably u_r-supported at $m \in M$*.

In words, the retailer's monetary profit from an unrandomized strategy $m \in M$ is favorably u_r-supported if, for all possible values of demand, the utility of this profit exceeds the utility of the threat strategy. Otherwise, it is unfavorably u_r-supported.

A nonlinear utility function invalidates the equal gains result implied by (A.1) under risk-neutrality. Instead, the payoff allocation in (13) follows the pattern in the next result which is a direct consequence of Theorem 4 in Roth and Rothblum (1982) (which generalizes Theorem 1 in Kihlstrom et al. (1981)).

Lemma 2. *Let $u_r(\cdot)$ and $\hat{u}_r(\cdot)$ be twice-differentiable and strictly increasing functions with respective Arrow-Pratt measures $R_r(\cdot)$ and $\hat{R}(\cdot)$ such that $R_r \lesssim \hat{R}$.*

(i) If the retailer's monetary payoff is favorably $\hat{u}_r$-supported at $m \in M$ for all contracts $m \in M$ that are given positive probability in (13) then the supplier's maximal expected utility $E\big[Eu_s(\pi_s(m))|i\big]$ is lower with u_r than with the more risk-averse $\hat{u}_r$.

(ii) If the retailer's payoff is unfavorably $\hat{u}_r$-supported at $m \in M$ for all contracts $m \in M$ that are given positive probability in (13) then the supplier's maximal expected utility $E\big[Eu_s(\pi_s(m))|i\big]$ is greater with u_r than with the more risk-averse $\hat{u}_r$.

Part (i) says that the supplier can achieve higher expected utility in (13) by replacing the u_r-retailer with the more risk-averse $\hat{u}_r$ if the $\hat{u}_r$-retailer's utility of monetary payoff exceeds the utility of the threat profit for all realizations of demand. Although this may seem rather technical, Lemma 3 (later in this section) identifies very simple conditions sufficient for this to occur. Finally, Theorem 1 obtains economic insight from Lemmas 2 and 3.

Part (ii) says that the supplier's maximal expected utility in (13) decreases when the u_r-retailer is replaced with the more risk-averse $\hat{u}_r$-retailer if there is a realization of demand for which the utility of profit of the $\hat{u_r}$-retailer is less than the utility from its threat strategy. In this case, the supplier's expected utility of profit in (13) *increases* as the retailer becomes *less* risk-averse. Although, again, the result seems rather technical, its application is straightforward: part (ii) of Lemma 2 applies when part (i) does *not* apply. Loosely speaking, this includes cases which are not addressed by Lemma 3.

Note that Lemma 2 discusses changes in the supplier's expected *utility* of profit in (13); it does not address changes in the expected profit itself. Since risk-return tradeoffs matter under risk-aversion, the proposition is a statement about first- and second-order stochastic dominance of the supplier's distribution of profit. Therefore, it is entirely possible in (13) for the supplier's maximal expected profit to increase while the supplier's maximal expected utility of profit, $b_s(i)$, decreases.

Example 1. As an example of this phenomenon, let $p = \$25, c_r = \$5, c_s = \$10, v = \$5,\ g_r = g_s = \$0$, with demand uniformly distributed on $[0, 100]$, and let $u_s(x) = u_r(x) = 1 - e^{-0.03x}$. The corresponding Arrow-Pratt coefficients of risk-aversion are $R_r(x) = R_s(x) = 0.03$ (for all x). If threat strategies yield utilities of zero to both players and the choice of supply contract types in (13) is reduced to the buy-back contract, then the solution of (13) yields a pair (expected utility, expected profit) for the supplier of $\{\$71.3401, 0.780285\}$.

In this example the retailer's payoff is unfavorably $\hat{u}_r$-supported if $u(\cdot)$ is replaced with $\hat{u}_r(x) = 1 - e^{-0.04x}$. Therefore, in accordance with Lemma 2, the supplier's expected utility is lower with the $\hat{u}_r$ retailer than with the u_r retailer. Specifically, it decreases from 0.780 to 0.776. At the same time, however, the supplier's expected profit increases from 71.3401 to 71.8443.

The next result is a consequence of axiom VI in Nash (1953) and can be regarded as a slightly weaker version of Corollary 2 which indicates that a player can increase its expected utility of profit in (13).

Remark 3. *In (13), player j's expected utility of profit increases as the expected utility from j's threat strategy payoff increases, $j \in \{r, s\}$.*

Hence, a risk-averse supplier has two basic modes to change its expected utility of profit in (13). Either it can seek a retailer with a different appetite for risk, or it can adopt a different threat strategy. The remainder of this section clarifies the risk preferences that a supplier should seek in a retailer.

4.3 Economic Insight From Revenue-Sharing and Buy-Back Contracts

As with risk-neutrality in §3, the Nash solution becomes simpler if the the set of feasible contracts is reduced to buy-back and revenue-sharing contracts with transfer payment functions given by (2a) and (2b). In particular, consider buy-back and revenue-sharing contract parameters, $\{w_b, b\}$ and $\{w_r, \phi\}$, that satisfy

$$(p - v + g_r - b) = \lambda(p - v + g) \qquad (w_b - b + c_r - v) = (\lambda - \theta)(c - v) \tag{14}$$

$$\phi(p - v) + g_r = \lambda(p - v + g) \qquad w_r + c_r - \phi v = (\lambda - \theta)(c - v) \tag{15}$$

where $0 \leq \lambda \leq 1$ and $\theta_l \leq \theta \leq \theta_u$. The size of the interval $[\theta_l, \theta_u]$ is implicit in the compactness of M which bounds the remaining parameters in (14) and (15). One can interpret (λ, θ) as parameters in a one-to-one mapping between the parameters (w_b, b) of the buy-back contract and the parameters (w_r, ϕ) of the revenue-sharing contract (cf. equations (2)).

It is interesting to compare the total supply chain profit, $\pi(q)$, given by (6), and the players' profits, $\pi_r(q, \lambda, \theta)$ and $\pi_s(q, \lambda, \theta)$, under the two types of contracts.

Proposition 7. *If a buy-back (revenue-sharing) contract has parameters that satisfy (14) and (15), then the players' profits are given by*

$$\pi_r(q, \lambda, \theta) = \lambda\pi(q) + \theta(c - v)q + D[\lambda g_s - (1 - \lambda)g_r] \tag{16a}$$

$$\pi_s(q, \lambda, \theta) = (1 - \lambda)\pi(q) - \theta(c - v)q - D[\lambda g_s - (1 - \lambda)g_r]. \tag{16b}$$

Clearly, Proposition 7 suggests a close connection between buy-back and revenue-sharing contracts. It is well-known that they are equivalent when the players are risk-neutral, and Proposition 7 shows that they *remain* equivalent even in the more general risk-averse setting.

Since (8) is a special case of (16) with $\theta = 0$, a comparison of the two pairs of equations yields an economic interpretation of (λ, θ). (Actually, it is possible to show that when $u_r(\cdot)$ and $u_s(\cdot)$ are *linear*, then without loss of optimality in (17), $\theta = 0$.)

As with risk-neutrality in §3, λ is the retailer's share of the total supply chain profit; the remainder goes to the supplier. Since $\pi(q)$ is a function of demand and D is a random variable, it follows that $\pi(q, D)$ is a random variable. So in (16), λ allocates the risk in the total supply chain profit $\pi(\cdot)$ between the retailer $\pi_r(\cdot)$ and the supplier $\pi_s(\cdot)$.

In contrast, since prices c and v are fixed and the order quantity, q, is deterministic, θ in

(16) allocates a risk-free transfer payment between the players. In other words, the inclusion of parameter θ permits a "redistribution" of risk between the supplier and the retailer. In particular, note that $\lambda = 0$ with $\theta > 0$ is a feasible allocation (as is $\lambda = 1$ with $\theta < 0$), which makes the retailer's (supplier's) profit completely risk free and obliges the supplier (retailer) to assume the entire supply chain risk.

Equations (16) sharpen the necessary and sufficient condition for a Nash bargaining solution (13) in a buy-back contract as follows:

$$\max_{\substack{\lambda\in[0,1]\\ \theta\in[\theta_l,\theta_u]\\ q\in m}} \Big[Eu_s\big(\pi_s(q,\lambda,\theta)\big) - Eu_s(\pi_s^{\tau_s})\Big]\Big[Eu_r\big(\pi_r(q,\lambda,\theta)\big) - Eu_r(\pi_r^{\tau_r})\Big], \tag{17}$$

where $\pi_s(\cdot)$ and $\pi_r(\cdot)$ are given by (16). Since $Eu_s[\pi_s(\cdot)]$ and $Eu_r[\pi_r(\cdot)]$ are continuous in q, λ and θ which take values in compact intervals, there exist q'^*, λ'^*, and θ'^* that are optimal in (17). Although their values cannot generally be given in closed form, q'^*, λ'^*, and θ'^* inherit the properties discussed earlier in this section. In particular, Corollary 3 implies that (i) $q'^* \leq q^*$, where q^* is the risk-neutral order quantity, and (ii) the players randomize arbitrarily between buy-back and the revenue-sharing contracts with parameters (w_b, b, w_r, ϕ) determined by (λ'^*, θ'^*) and equations (14) and (15).

The following sequel to Lemma 2 gives conditions for the supplier's expected utility to increase as the retailer becomes more risk-averse. The result is presented in the context of buy-back and revenue-sharing contracts and transformations (14) and (15).

Suppose in (17) that the supplier (retailer) considers a succession of replacements of its retailer (supplier). Let $R_1(x), R_2(x), \ldots$ be the corresponding sequence of the retailers' (suppliers') coefficients of risk-aversion. It is convenient to use the notation $R \rightsquigarrow 0$ for the convergence of $R_1(x), R_2(x), \ldots$ to 0 for all x, and $R \rightsquigarrow \infty$ for the divergence of $R_1(x), R_2(x), \ldots$ to ∞ for all x. Let $R_s(\cdot)$ and $R_r(\cdot)$ be the coefficients of risk-aversion that correspond to $u_s(\cdot)$ and $u_r(\cdot)$, respectively.

Lemma 3. *Assume $g_s = g_r = 0$. Let q'^*, λ'^*, and θ'^* be optimal in (17), and let $q_s'^*$ denote a supply quantity, q, that maximizes $Eu_s(\pi(q))$, where $\pi(q)$ is the total supply chain profit given by (6). If (i) $R_s \rightsquigarrow 0$ or (ii) $R_r \rightsquigarrow \infty$, then $q'^* \to q_s'^*$ and $\lambda'^* \to 0$.*

Using part (ii) of Lemma 3 and equation (16a), one can show that there exists a coefficient of risk-aversion, $\bar{R}(\cdot)$ say, such that, for all retailer's coefficients of risk-aversion $R_r(\cdot)$ with $R_r \gtrsim \bar{R}$, the retailer's profit becomes essentially deterministic. To gain some economic intuition, note that part (ii) implies that as $R_r(\cdot)$ "becomes very large", λ optimal in (17) decreases, causing the retailer's utility of profit to become $u_r[\theta(c - v)q] \geq 0$. Then from Definition 3, the retailer's monetary profit is favorably u_r-supported. We refer to a retailer with $R_r \gtrsim \bar{R}$ as *very risk-averse*. Part (i) of Lemma 3 reveals that the same phenomenon occurs when the supplier is risk-neutral.

The following statement summarizes this discussion of the consequences of Lemmas 2 and 3.

Theorem 1. *Suppose $g_s = g_r = 0$. A supplier can increase its expected utility of profit in a supply chain by using buy-back or revenue-sharing contracts and by choosing a retailer that is*

(i) either risk-neutral or very risk-averse if the supplier is risk-averse

(ii) risk-averse if the supplier is risk-neutral.

Theorem 1 reveals that a supplier has an incentive to choose its retailer based on the retailer's risk preferences. Part (i) stipulates that a risk-averse supplier favors retailers with extreme risk preferences – that is, retailers that are either risk-neutral (or, equivalently, least averse), or retailers that are most risk-averse. Therefore, if given a choice, a risk-averse supplier should choose the most or the least risk-averse "candidate in the room." Choosing any other retailer would diminished the supplier's expected utility of profit.

Both types of extremes benefit the risk-averse supplier in very different ways. A presence of the risk-neutral retailer raises the optimal order quantity, and eliminates the burden of demand risk for the risk-averse supplier. In fact, the risk-neutral retailer, bears the entire supply chain risk (to prove this, in Lemma 3, simply swap the names given to the players).

The exact opposite happens when the risk-averse supplier enters in a contract with a very risk-averse retailer. Then it is the risk-averse supplier absorbing the supply chain risk in (17) while extracting most of the retailer's rent in the process. Although the optimal order quantity in (17) decreases (cf. Proposition 5), the additional rent extracted from the retailer more than offsets the loss of the supply chain performance.

Part (ii) of Theorem 1 reveals that the choice of a retailer is even simpler when the supplier is risk-neutral. Given a choice, a risk-neutral supplier should always choose a risk-averse retailer over a risk-neutral one.

Example 2. Here, we reconsider the example that follows Lemma 2 and demonstrate (ii) of Theorem 1. We already know that the supplier's expected utility decreases when the u_r-retailer with $u_r(x) = 1 - e^{-0.03x}$ is replaced with $\hat{u}_r = 1 - e^{-0.04x}$. Now consider replacement with a different $\hat{u}_r$-retailer with $\hat{u}_r = 1 - e^{-0.3x}$. Solving (17) yields $\lambda'^* \approx 0$, and $\theta'^* = 0.12$. The retailer's payoff is now favorably $\hat{u}_r$-supported (inspect (16a)) because it is independent of the stochastic demand, D. Consistent with Lemma 2, the supplier's pair (expected profit, expected utility) increases from ($71.3401, 0.780285) to ($81.7225, 0.815667) while the total expected supply chain profit decreases from $142.672 to $96.9654.

In this example it can be shown that a *risk-neutral* retailer's expected utility of profit is unfavorably $\hat{u}_r$-supported for all buy-back contracts. So Lemma 2 (ii) implies that the supplier can also increase its expected utility of profit by replacing the u_r-retailer ($u_r(x) = 1 - e^{-0.03x}$) with one which is risk-neutral. The pair (expected profit, expected utility of profit) that the supplier attains from doing so is ($72.0412, 0.85). That is, the supplier's expected profit is higher ($81.7225 versus $72.0412) when the u_r-retailer is replaced with the more risk-averse one (in comparison to the risk-neutral one). The expected utilities of profit, 0.815667 versus 0.85, are slightly in favor of replacing the u_r-retailer with the risk-neutral one.

5. Example

This section uses an example to contrast the risk-neutral and risk-averse cases. As in Example 1, let $p = \$25, c_r = \$5, c_s = \$10, v = \$5$, $g_r = g_s = \$0$, and let demand be uniformly distributed on $[0, 100]$. Also, assume $Eu_s(\pi_s^{\tau_s}) = Eu_r(\pi_r^{\tau_r}) = 0$. If the players are risk-neutral, then (7) and (11) imply that the risk-neutral order quantity in the Nash bargaining solution is $q^* = 100(\$25 - \$15)/(\$25 - \$5) = 50$ and $\lambda^* = 1/2$. That is, the players should evenly split the total supply chain profit induced by the Nash-optimal order quantity of 50 units. This can be achieved, for example, with a buy-back contract having parameters $w_b = \$15$ and $b = \$10$. So the supplier charges the retailer a wholesale price of $15 and repurchases unsold inventory at the end of the selling season at $10 per unit.

Suppose instead that the players are risk-averse with $u_r(x) = x - \alpha x^2$ and $u_s(x) = x - \beta x^2$ where $\alpha = 0.001$, and $\beta = 0.003$. The Nash-optimal values from (17) are $q'^* = 33, \lambda'^* = 0.82$, and $\theta'^* = -0.14$. Risk-aversion diminishes the Nash-optimal order quantity from 50 units to 33 units. Using λ'^*, θ'^*, (14), and (15) to determine the Nash-optimal contract parameters (and hence the division of the monetary payoffs), $w_b = \$13.21, b = \$3.60$, $w_r = \$7.30$, and $\phi = \$0.82$. That is, under a buy-back contract the supplier charges the retailer $13.21 per unit and repurchases unsold inventory at $3.60 per unit. Under a revenue-sharing contract, the supplier charges the retailer $7.30 per unit, the retailer retains 82% of the sales revenue, and passes the remaining 18% to the supplier.

If the retailer is risk-averse but the supplier is risk-neutral ($\beta = 0$ and $u_s(x) = x$ for all x), (17) implies $q'^* = 50, \lambda'^* = 0$ and $\theta'^* = 0.23$. This confirms the result in Lemma 3 and it implies that the risk-neutral supplier assumes the entire supply chain risk. To see this, the parameters $\lambda'^{*} = 0$ and $\theta'^{*} = 0.23$ and (16) yield the following monetary profits:

$$\pi_r(q) = 0.23(c - v)q \qquad\qquad \pi_s(q) = \pi(q) - 0.23(c - v)q.$$

So the retailer's profit becomes deterministic and depends only on the order quantity, q, while the supplier bears the risk embedded in $\pi(q)$ given by (6). The buy-back and revenue-sharing contracts that achieve these profits have parameters $w_b = \$17.70, b = \$20$, $w_r = \$ - 7.40$, and $\phi = 0$.

The revenue-sharing contract with parameters $w_r = \$-7.40$ and $\phi = 0$ has an interesting interpretation because the retailer receives none of the supply chain revenue, but the supplier compensates the retailer $\$7.40q$ for placing an order of size q. Therefore the supplier merely "rents" shelf space from the retailer.

6. Generalizations: Price-Dependent Demand

The supplier, as well as the retailer, has a financial interest in the retail price. The probability distribution of demand depends implicitly, on the retail price, and a lower price would likely induce a stochastically greater demand. Hence, an appropriate order quantity (and the supplier's revenue from the outbound shipment) would be greater. Thus, there is an incentive for the retail price to be included in the bargaining process. In this section we treat the retail price endogenously, constrain it to an interval, say $[p_l, p_u]$, and let $F_p(\cdot)$ denote the distribution function of demand when the retail price is p. We assume that the price is set at the same time as the remaining parameters in the model, namely the vector $m \in M$. Although this assumption is commonly made in the research literature on supply chain contracts, it precludes a dynamic pricing strategy that adjusts the retail price during the selling season. We note that this description has not yet specified the player(s) who sets the price.

Although pricing decisions typically rest with retailers, there are some products, such as books and automobiles, whose suppliers influence retail prices via MSRPs (manufacturer's suggested retail price). That is, the suppliers in some large industries influence retail prices. So the first subsection considers the extreme case in which pricing decision is made jointly by the retailer and supplier. There, the players *bargain* over the retail price. Then we consider the case where the retailer alone sets the retail price but the players anticipate the dependence of the price-setting decision on the choice of $m \in M$. The supplier is worse off in the latter case because the retailer may choose a different price than the one that would have emerged from the bargaining process (in the former case).

The model remains otherwise unchanged from §2. For ease of exposition, we assume that the retailer and supplier are risk-neutral and restrict the supply contract types to buy-back and revenue-sharing with transfer payment functions given by (2a) (2b).

6.1 Joint Pricing Decision

Since the retailer and supplier prefer different prices (cf. §6.2), here we let p be settled by way of negotiation. The principal results are that Propositions 1, 2, Corollary 2, and Remark 1 remain valid with the caveat that the order quantity and contract parameters become functions of p. Such supply contracts are closely related to the price-contingent (or "bill-back") contracts studied by Bernstein and Federgruen (2005).

In contrast to §2, we augment $m \in M$ with the retail price p. Formally, replace m with $m' = (p, m)$ and M with $M' = [p_l, p_u] \times M$ which remains a non-empty and compact set of pure strategies. In particular, m' is a vector that consists of the retail price, p, and a vector $m \in M$. Let $\Delta(M')$ denote the set of all probability distributions over the set of Borel measurable subsets of M' and, in (3), replace m and M with m' and M' respectively.

If the retail price is exogenous, the Nash bargaining solution obliges the players to randomize at most in their choice of the type of supply contract (cf. §3). This conclusion stems from the concavity of $\pi_r(\cdot)$ and $\pi_s(\cdot)$ on M (cf. (1)) for each outcome of demand, and Lemmas 1 and 4. To obtain the same result here, it is sufficient for the players' profit functions to be concave functions of $(p, m) \in M'$ for each outcome of demand. The restriction to buy-back and revenue-sharing contracts reduces the needed assumption to concavity of the integrated channel profit function $\pi(q, p)$ as a function of p.

The steps that lead to (9) yield the following necessary and sufficient condition for a Nash-bargaining solution:

$$\max_{\substack{\lambda\in[0,1]\\ q,p\in m}} \big[(1-\lambda)E\pi(q,p) - \mu\lambda g + \mu g_r - E\pi_s^{\tau_s}\big]\big[\lambda E\pi(q,p) + \mu\lambda g - \mu g_r - E\pi_r^{\tau_r}\big] \tag{18}$$

and

$$q^*(p) = F_p^{-1}\left(\frac{p-c+g}{p-v+g}\right) \tag{19a}$$

$$\lambda^*(p) = \frac{E\pi(q,p) + E\pi_r^{\tau_r} - E\pi_s^{\tau_s} + 2\mu g_r}{2\big(E\pi(q,p) + \mu g\big)} \tag{19b}$$

Here, $q^*(p)$ and $\lambda^*(p)$, the order quantity and the retailer's fraction of the total supply chain profit, are functions of the retail price, p.

The existence of p^* is guaranteed because p takes values on a compact interval and the maximand in (18) is continuous in p. Concavity of $E\pi(q, p)$ with respect to p and the first-order condition with respect to p of (18) yield:

$$\frac{\partial E\pi(q,p)}{\partial p}\Big\{(1-\lambda)\big[\lambda E\pi(q,p)+\mu\lambda g-\mu g_r-E\pi_r^{\tau_r}\big] \\ +\lambda\big[(1-\lambda)E\pi(q,p)-\mu\lambda g+\mu g_r-E\pi_s^{\tau_s}\big]\Big\}=0, \qquad (20)$$

where (q,λ) is given by the right side of (19).

Therefore, the supplier and the retailer choose a quantity-price pair, (q^*,p^*), that maximizes the expected total profit of the supply chain. The players opt for a buy-back or revenue-sharing contract with parameters chosen so that the retailer and the supplier receive fractions λ^* and $1-\lambda^*$ of the expected total supply chain profit.

6.2 Unilateral Pricing Decision

This section considers the scenario where the retailer *alone* sets the retail price, p, and this significantly changes the Nash solution in §6.1. First, the order quantity and the expected profit allocation now depend on the contract. Second, the supplier and the retailer are no longer indifferent among the types of supply contracts; this leads to randomized strategies. Using the revenue-sharing contract as an example, we show how the pricing decision allows the retailer to extract additional expected profit from the supplier. Hence, the pricing decision plays a role that is similar to risk-aversion (cf. §4), where there is a redistribution of expected profit due to different attitudes toward risk, and that favors the less risk-averse player.

The model can be analyzed by assuming that the game between the supplier and the retailer remains unchanged from that described in §2, except that the *exogenous* retail price, p, now becomes *endogenous* and is controlled by the retailer who seeks to maximize expected profit, namely $E\pi_r(p,m)$. This introduces an additional optimization constraint on the solution of the bargaining game between the supplier and the retailer which can be written as:

$$\max\big[E(E\pi_s(m)|i)-E\pi_s^{\tau_s}\big]\big[E(E\pi_r(m)|i)-E\pi_r^{\tau_r}\big] \qquad (21a)$$

$$\text{s.t.} \quad \frac{\partial E\pi_r(p,m)}{\partial p}=0,\ \forall\ i \qquad (21b)$$

The maximization in (21a) is with respect to $m\in M$ and with respect to probabilities of choosing a particular supply contract type. The constraint (21b) holds for the two i corresponding to buy-back and revenue-sharing contracts.

Because the retail price, p, optimal in (21b) depends on the contract type and the order quantity, $q(p)$, optimal in (21a) is a function of the retail price, p, the following result follows:

Proposition 8. *Let the profit-maximizing retailer set the retail price, p. Then the order quantity, q, is contract type specific.*

Since the expected profit earned by retailer and the supplier now also becomes contract specific, the players are no longer indifferent among the types of supply contracts. Moreover, the contract most preferred by the retailer becomes the least preferred type by supplier and vice versa. This yields the following result:

Corollary 4. *Let the profit-maximizing retailer set the retail price, p. Then, in (21), the players can be assumed to randomize in their choice of supply order quantity and supply contract type.*

When the choice of supply contracts is reduced to a single type, say revenue-sharing, then the retailer's expected profit is given by:

$$E\pi_r(q, m) = (\phi(p - v) + g_r))S(q, p) - (w_r + c_r - \phi v)q - g_r\mu, \tag{22}$$

obtained by taking the expectation of (1a).

The solution is then a non-randomized strategy given by (19) and a retail price, p^*, which maximizes (22). It is well-known that the parameter λ in (19) and the revenue-sharing contract parameters (ϕ, w_r) in (22) can be mapped as follows:

$$\phi(p - v) + g_r = \lambda(p - v + g) \qquad \text{and} \qquad w_r + c_r - \phi v = \lambda(c - v)$$

Intuitively, when compared to the case analyzed in §3, the pricing decision should afford the retailer some additional bargaining power. From the results in §3, the retailer's and the supplier's ability to secure a better deal in the bargaining game rests with the expected payoffs of their respective threat strategies. Here, the retailer has the ability to capture a larger portion of the total expected supply chain profit through the pricing decision as well. To see this, note that for the case of $g_r = g_s = 0$, the retail price, p^*, that maximizes $E\pi_r(m)$ also maximizes $\lambda^*(p)$ as a function of p (cf. 19b).

We conclude by noting that using the traditional equilibrium point approach, it is sometimes not all obvious how to design an equilibrium contract (cf. Cachon and Lariviere (2005) for equilibrium contract design using revenue-sharing). This is because the incentives to coordinate the retailer's quantity decision frequently distort the retailer's price decision. Therefore solving the bargaining game offers an alternative way to proceed. The analytical elegance from §3, however, is lost.

7. Conclusion

The research literature on supply chain contracts commonly assumes implicitly that firms employ risk-neutral evaluations of market risk; so they replace random monetary payoffs with expected values. Frequently, the supply chain is modeled as a Stackelberg game with a specific type of contract and one of the firms is designated as the "leader." The other firm has no recourse but to accept its leadership, so there is a unique division of the total supply chain's profit in which the leader expropriates the follower's profit until the follower earns no more than its reservation level. Therefore, the Stackelberg game endows the leader with much more bargaining power than is realistic in many supply chains.

In this paper, we consider the case in which *both* firms have bargaining power and are sensitive to risk; this includes the case of two risk-neutral firms with bargaining power. In this setting, we study how the firms would choose among alternative types of supply contracts and the worths of their bargaining positions. Thus, we formulate a pure bargaining game between the supplier and retailer and solve it using the Nash bargaining solution.

We show that risk aversion has a significant impact on a bargaining position and on the value of a potential supply chain contract. It is well-known that risk-aversion adversely affects the performance of a supply chain by reducing the optimal order quantity, and that lowers the supply chain's total expected profit. We confirm this by comparing the Nash-optimal order quantities in the risk-neutral and risk-averse cases.

The solution to the bargaining problem also reveals a particularly important strategic insight. In spite of the reduced expected profit, one player's risk-aversion may be advantageous to the other. To see this, under risk-neutrality the Nash bargaining solution awards both players equal gains in the expected profit over their expected reservation payoffs. However, if one of the players is risk neutral, we prove that the expected profit of that player increases if the other player is more risk-averse. Therefore *risk-neutral* suppliers (retailers) have an incentive to choose *risk-averse* retailers (suppliers) as their supply chain partners. We further prove that *risk-averse* suppliers (retailers) attain the highest expected utility of profit by choosing the most or the least risk-averse retailers (suppliers) available. These findings run counter to the existing literature on supply chain contracting, which generally focuses on risk intermediation or hedging. We conclude the paper by finding an endogenous Nash-optimal retail price for the product.

The analysis is predicated on some limiting assumptions. Nash (1950, 1953) assumes that there is no private information. However, Harsanyi and Selten (1972) show that a solution of a two-person bargaining problem with incomplete information reduces to a solution of a generalized Nash bargaining game. Neslin and Greenhalgh (1983) present experimental evidence that supports the ability of Nash's theory to predict outcomes of buyer-seller negotiations. Nagarajan and Bassok (2008) discuss this point and the application of insights from Nash's bargaining model to supply chain negotiations. It may be worth exploring this issue in future research.

Some studies also use a "*generalized* Nash bargaining solution (GNB)" (e.g. Nagarajan and Bassok (2008)). See Roth (1979) for a detailed description of GNB. A comparison of the results presented here and those obtained via GNB, however, reveals that the latter merely introduces an additional scalar factor into some of the results, but preserves their qualitative nature.

Appendix

A. Proofs of Statements

Proof. PROOF OF PROPOSITION 1. The first-order conditions with respect to contract parameters reduce to:

$$E(E\pi_s(m)|i) - E\pi_s^{\tau_s} = E(E\pi_r(m)|i) - E\pi_r^{\tau_r}. \tag{A.1}$$

□

Proof. PROOF OF COROLLARY 1. Let $a = E(E\pi_s(m)|i) - E\pi_s^{\tau_s}$ and $b = E(E\pi_r(m)|i) - E\pi_r^{\tau_r}$. In (5), $m \in M$ and the probabilities of choosing various types of supply contracts maximize ab, or equivalently, $\sqrt{ab}$. Since, by Proposition 1, $a = b$, then $\sqrt{ab} = (1/2)(a + b)$. □

Proof. PROOF OF PROPOSITION 2. Use (A.1) to reduce the following first-order condition with respect to q

$$\frac{E(E\pi_s(m)|i)}{\partial q}\left[E(E\pi_s(m)|i) - E\pi_s^{\tau_s}\right] + \frac{E(E\pi_r(m)|i)}{\partial q}\left[E(E\pi_r(m)|i) - E\pi_r^{\tau_r}\right] = 0. \tag{A.2}$$

to

$$\frac{E(E\pi_s(m)|i)}{\partial q} + \frac{E(E\pi_r(m)|i)}{\partial q} = \frac{E\pi(q)}{\partial q} = 0, \tag{A.3}$$

which holds at $q = q^*$. □

Proof. PROOF OF PROPOSITION 5. Let $h(\cdot)$ denote the maximand in (13). Construct $\hat{h}(\cdot)$ by replacing $u_r(\cdot)$ in $h(\cdot)$ with $\hat{u}_r(\cdot)$, where the Arrow-Pratt measures of risk-aversion that correspond to $u_r(\cdot)$ and $\hat{u}_r(\cdot)$ satisfy $R_r \lesssim \hat{R}_r$. Let $m \in M$ and $\hat{m} \in M$ be optimal in $\max(h(m))$ and $\max(\hat{h}(\hat{m}))$ respectively. Without loss of optimality, the vector i is fixed; that is, the choice of a contract type is fixed (cf. Remark 4).

Using the notation adopted in §2, the first component the $m \in M$, is the supply quantity q. Similarly, the first component of $\hat{m} \in M$, is the supply quantity $\hat{q}$. By contradiction, suppose $q \leq \hat{q}$.

We may scale $u_r(\cdot)$ in $h(\cdot)$ and $\hat{u}_r(\cdot)$ in $\hat{h}(\cdot)$ so that

$$h(m) = \hat{h}(m). \tag{A.4}$$

By optimality of $\hat{m} \in M$ we have

$$\hat{h}(\hat{m}) \geq \hat{h}(m). \tag{A.5}$$

Properties of functions with differing degree of risk-aversion imply

$$h(\hat{m}) \geq \hat{h}(\hat{m}). \tag{A.6}$$

Now combine (A.4), (A.5), and (A.6) to obtain $h(\hat{m}) \geq h(m)$. However, since $m \in M$ is optimal in $\max(h(m))$, then we also have $h(\hat{m}) \leq h(m)$, a contradiction. $\square$

Proof. PROOF OF PROPOSITION 6. Let $\pi(q)$ be the total supply chain profit given by (6) and let the order quantities q^* and q'^* be optimal in (5) and (13), respectively. From (A.3), $\partial E\pi(q^*)/\partial q = 0$ and by concavity $\partial E\pi(q)/\partial q \geq 0$ for all $0 \leq q \leq q^*$. The result follows from Corollary 3, $q'^* \leq q^*$, and Proposition 5. $\square$

Proof. PROOF OF PROPOSITION 7. We prove (16a) only for a buy-back contract; the proof is similar for a revenue-sharing contract. Using (1a), (6), and (14), the retailer's profit from a buy-back contract is

$$\begin{aligned}
\pi_r &= (p - v + g_r - b)S - (w_b - b + c_r - v)q - g_r D \\
&= \lambda(p - v + g)S - (\lambda - \theta)(c - v)q - g_r D \\
&= \lambda[(p - v + g)S - (c - v)q - gD] + \theta(c - v)q + \lambda g D - g_r D \\
&= \lambda\pi(q) + \theta(c - v)q + \lambda g D - g_r D,
\end{aligned}$$

which is (16a). Expression (16b) follows from substitution for $\pi(q)$ and $\pi_r(q,\lambda,\theta)$ in $\pi(q) = \pi_r(q,\lambda,\theta) + \pi_s(q,\lambda,\theta)$ and simple algebra. $\square$

Proof. PROOF OF LEMMA 1. Lemma 4 in Appendix B establishes that for a given supply contract type, the supplier and the retailer need not randomize to choose the order quantity, q, and contract parameters. Concavity of $\pi_s(m), \pi_r(m)$ in $m \in M$ (for each outcome in the sample space of D), Jensen's inequality, and an argument that is analogous to that used in the proof of Lemma 4 reveal that both players choose the same order quantity, q, for all i, i.e. for all contract types. $\square$

Proof. PROOF OF LEMMA 3. Only (i) is proved; the proof of (ii) is analogous. To streamline the proof, suppose $u_s(\pi_s^{\tau_s}) = u_r(\pi_r^{\tau_r}) = 0$. Let $\pi_s(q,\lambda,\theta)$ and $\pi_r(q,\lambda,\theta)$ be given by (16). As $R_s \rightsquigarrow 0$, $q = q_s'^*$ maximizes the expected profit of the entire supply chain, $E\pi(q)$. Continuity then yields the existence of $\theta'^* > 0$ such that

$$[E\pi(q'^*) - \theta'^*(c - v)q'^*]u_r\big(\theta'^*(c - v)q'^*\big) \geq [E\pi_s(q,\lambda,\theta)]Eu_r\big(\pi_r(q,\lambda,\theta)\big)$$

for all feasible q, λ, and θ. $\square$

B. Lemmas

Lemma 4. *Let the sets U and U' be defined as follows:*

$$U' = \{(Eu_s[\pi_s(m)], Eu_r[\pi_r(m)]) : m \in M\} \quad \text{and} \quad U = \{(Eu_s(\pi_s^\gamma), Eu_r(\pi_r^\gamma)) : \gamma \in \Delta(M)\},$$

where $u_j(\cdot)$, $j \in \{r, s\}$ is concave increasing, $\pi_j(\cdot)$, $j \in \{r, s\}$ is concave on M for each outcome in the sample space of D, and the vector i is fixed; that is, the choice of a contract type is fixed. Then $U = U'$.

Proof. PROOF OF LEMMA 4. Clearly $U' \subseteq U$. Since $u_j(\cdot)$ is concave increasing, then $u_j(\pi_j(m, D))$, and consequently $Eu_j(\pi_j(m, D))$, are concave in $m \in M$ for each outcome in the sample space of D. The latter follows because linear combinations of concave functions are concave. Concavity and Jensen's inequality then yield

$$Eu_j[\pi_j(Em, D)] \geq E\big(Eu_j[\pi_j(m, D)]|m\big) = Eu_j(\pi_j^\gamma).$$

Now letting $m_3 = Em$, the above inequality can written as

$$Eu_j(\pi_j(m_3, D)) \geq Eu_j(\pi_j^\gamma).$$

Note that m_3 is a pure strategy, and with i fixed, $m_3 \in M$. Therefore we found a pure strategy that is at least as preferable to both players as any mixed strategy that for a given supply contract type randomizes to choose an order quantity and contract parameters. Then $U \subseteq U'$ and $U = U'$. $\square$

Acknowledgements

The authors gratefully acknowledge comments by James Alexander, Lingxiu Dong, and Tava Lennon Olsen which improved the paper.

References

Agrawal, V., S. Seshadri. 2000. Risk intermediation in supply chains. *IIE Transactions* **32** 819–831.

Bernstein, F., A. Federgruen. 2005. Decentralized supply chains with competing retailers under demand uncertainty. *Management Sci.* **51**(1) 18–29.

Cachon, G. 2003. Supply chain coordination with contracts. *Handbooks in Operations Research and Management Science, 11: Supply Chain Management: Design, Coordination and Operation.* Chapter 6 of Graves, S. and de Kok, T. eds., North-Holland.

Cachon, G., S. Netessine. 2004. Game theory in supply chain analysis. *Handbook of Quantitative Supply Chain Analysis: Modeling in the E-Business Era* Chapter 2 of Simchi-Levi, D., Wu, S. D. and Shen, M. eds, Kluwer.

Cachon, G. P., M. A. Lariviere. 2005. Supply chain coordination with revenue-sharing contracts: Strengths and limitations. *Management Sci.* **51**(1) 30–44.

Chod, J., N. Rudi. 2006. Strategic investments, trading, and pricing under forecast updating. *Management Sci.* **52**(12) 1913–1929.

Dong, L., H. Liu. 2005. Equilibrium forward contracts on non-storable commodities in the presence of market power. *Oper. Res.* **55**(1) 128–145.

Gan, X., S. P. Sethi, H. Yan. 2004. Coordination of supply chains with risk-averse agents. *POMS* **13**(2) 135–149.

Gurnani, H., M. Shi. 2006. A bargaining model for a first-time interaction. *Management Sci.* **52**(6) 865–880.

Harsanyi, J. C., R. Selten. 1972. A generalized Nash solution for two-person bargaining games with incomplete information. *Management Sci.* **18**(5) 80–106.

Kihlstrom, R. E., A. E. Roth, D. Schmeidler. 1981. Risk aversion and solutions to Nash's bargaining problem. *Game Theory and Mathematical Economics: Proceedings of the Seminar on Game Theory and Mathematical Economics Bonn/Hagen, 7-10 October, 1980* Moeschlin, O., Pallaschke, D. eds, North-Holland Publishing Company.

Kohli, R., H. Park. 1989. A cooperative game theory model of quantity discounts. *Management Sci.* **35**(6) 693–707.

Myerson, R. B. 2004. *Game Theory: Analysis of Conflict*. 6th ed. Harvard University Press, Cambridge, MA, 02138.

Nagarajan, M., Y. Bassok. 2008. A bargaining framework in supply chains: The assembly problem. *Management Sci.* **52**(12) 1913–1929.

Nagarajan, M., G. Sošić. 2008. Game-theoretic analysis of cooperation among supply chain agents: Review and extensions. *Eur. J. Oper. Res.* **187** 719–745.

Nash, J. 1950. The bargaining problem. *Econometrica* **18**(2) 155–162.

Nash, J. 1953. Two-person cooperative games. *Econometrica* **21**(1) 128–140.

Neslin, S. A., L. Greenhalgh. 1983. Nash's theory of cooperative games as a predictor of the outcomes of buyer-seller negotiations: An experiment in media purchasing. *Journal of Marketing Research* **20**(4) 368–379.

Osborne, M. J., A. Rubinstein. 1994. *A Course in Game Theory*. MIT Press, Cambridge, MA, 02142-1493.

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

Pindyck, R., D. Rubinfeld. 2008. *Microeconometrics*. 7th ed. Prentice Hall, Upper Saddle River, NJ, 07458.

Reyniers, D. J., C. S. Tapiero. 1995. The delivery and control of quality in supplier-producer contracts. *Management Sci.* **41**(10) 1581–1589.

Roth, A. E. 1979. *Axiomatic Models of Bargaining*. Springer-Verlag, New York, NY 10010.

Roth, A. E., U. G. Rothblum. 1982. Risk aversion and Nash's solution for bargaining games with risky outcomes. *Econometrica* **50**(3) 639–647.

Rubinstein, A. 1982. Perfect equilibrium in bargaining model. *Econometrica* **50**(1) 97–109.

Shubik, M. 1982. *Game Theory in the Social Sciences: Concepts and Solutions*. MIT Press, Cambridge, MA, 02142-1493.

The Wall Street Journal. 2006. Enterprise: Working things out with a giant customer (October 17, 2006) B4.

Tomlin, B. 2003. Capacity investments in supply chains: Sharing the gain rather than sharing the pain. *Manufacturing Service Oper. Management* **5**(4) 317–333.

Van Mieghem, J. A. 1999. Coordinating investment, production, and subcontracting. *Management Sci.* **45**(7) 954–971.

Van Mieghem, J. A. 2003. Capacity management, investment, and hedging: Review and recent developments. *Manufacturing Service Oper. Management* **5**(4) 269–302.

Williams, K. 2008. Supply chain risk: Are you prepared? *Strategic Finance* **89**(11) 23–24.

Wu, S. D. 2004. Supply chain intermediation: A bargaining theoretic framework. *Handbook of Quantitative Supply Chain Analysis: Modeling in the E-Business Era* Chapter 3 of Simchi-Levi, D., Wu, S. D. and Shen, M. eds, Kluwer.

Zeuthen, F. 1930. *Problems of Monopoly and Economic Welfare*. G. Routledge and Sons, London.