Shared-Savings Contracts for Indirect Materials in Supply Chains: Channel Profits and Environmental Impacts

Charles J. Corbett • Gregory A. DeCroix
The Anderson School at UCLA, 110 Westwood Plaza, Box 951481, Los Angeles, California 90095-1481
The Fuqua School of Business, Duke University, Durham, North Carolina 27708
charles.corbett@anderson.ucla.edu • decroix@mail.duke.edu

There are many materials for which the quantity needed by a firm is at best indirectly related to the quantity of final product produced by that firm, such as solvents in manufacturing processes or office supplies. For any such "indirect" materials, an inescapable incentive conflict exists: The buyer wishes to minimize consumption of these indirect materials, while the supplier's profits depend on increasing volume. Both buyer and supplier can exert effort to reduce consumption, hence making the overall supply chain more efficient. However, no supplier will voluntarily participate unless contract terms are fundamentally revised. This can be done through a variety of "shared-savings" contracts, where both parties profit from a consumption reduction. This paper analyzes several such contracts currently in use for chemicals purchasing. We show that such contracts can always increase supply-chain profits but need not lead to reduced consumption. We analyze equilibrium effort levels, consumption, and total profits, and show how these change with the contract parameters. We find that the goals of maximizing joint profits and minimizing consumption are generally not aligned. Also, surprisingly, a decrease in a cost parameter can lead to a decrease in profits; it may be necessary (but is always possible) to renegotiate the shared-savings contract to reap the benefits of a cost decrease.
(*Supply-Chain Management; Supply Contracts; Shared Savings; Game Theory; Environmental Management; Indirect Materials*)

1. Introduction

The potential for improved supply-chain performance through various forms of coordination has been demonstrated extensively in recent literature and management practice. A wide range of supply contracts has been studied, generally with the intention of reducing total supply-chain inventory costs. However, this literature (implicitly) assumes that the materials in question will become part of the product sold to final customers. This means that demand for the materials is directly linked to final demand, assumed to be constant or price-sensitive through a (possibly stochastic) demand curve. An automobile always requires four wheels, and an assembler's usage of wheels is directly linked to final car sales. Many materials, however, are *indirect*: They are consumed during the production process but do not become part of the final product. Demand for such indirect materials is, by definition, only indirectly linked to final market demand; solvents are needed at various steps in the automotive assembly process, but consumption of these solvents depends as much on efficiency of usage as on demand for cars. Even consumption of materials such as paint, which does become part of the final car, can often be reduced by less wasteful spraying processes. In this paper, we

0025-1909/01/4707/0881$5.00
1526-5501 electronic ISSN

examine the impact of supply contracts on efforts to reduce consumption of such indirect materials.

Reduction of consumption of indirect materials entails a cost reduction for the customer. For the supplier, however, this means a loss of volume, hence almost certainly lower profits. This simple observation lies at the heart of a major incentive conflict inherent in almost any supply chain: Suppliers can earn more by selling more, even if the material is undesirable from a broader perspective. This is acutely true in the case of environmentally undesirable products, such as hazardous chemicals or CFC-based solvents. Manufacturers may need these to keep their processes running, but would much rather consume less or find a way of doing without altogether. However, they will be hard-pressed to convince their suppliers to help them achieve this without offering some form of compensation.

Recently, new types of chemical contracts have emerged, under which the basis of payment shifts from quantity to service (Bierma and Waterstraat 1996, 2000), also referred to as "servicizing" (Reiskin et al. 2000). These include chemical management fees and lease arrangements (where payment is based partly on quantity and partly on additional services, and the supplier retains ownership of the chemicals). Shared-savings contracts go further, explicitly measuring the benefits accruing to both parties from costs avoided by reducing consumption, and sharing these. The underlying belief is that these contracts will always increase profits and reduce consumption compared to the original quantity-based model.

In this paper, we evaluate these different types of contracts from a theoretical and practical perspective. For any given contract, we characterize the level of effort that supplier and customer will exert in equilibrium to reduce consumption, and compare the resulting consumption and profit levels under various contracts. Though there may be multiple equilibria, we find that the set of equilibria is a lattice. Contrary to commonly held beliefs, shared-savings contracts will not always reduce consumption, and the goals of minimizing total consumption and maximizing joint profits are generally not aligned. However, we show that it is always possible to increase channel profits with a simple shared-savings contract. While we show that shared-savings contracts cannot achieve full channel coordination, our numerical experiments suggest that an appropriate shared-savings contract tends to deliver a high percentage of the potential gain that channel coordination could achieve over a traditional quantity-based contract.

As the unit price decreases, the supplier will always exert more effort, the customer less. (This is a stronger result than the traditional comparative statics on a lattice, which only states that the entire lattice of equilibria will shift in a particular direction, and does not predict how individual nonextreme points will change.) Similarly, if the supplier's manufacturing costs increase, he will exert more effort and the customer less, while an increase in customer disposal costs has the opposite effect. Surprisingly, though, a decrease in either cost could lead to a *decrease* in total profits, unless the cost decrease is shared between both parties by adjusting the unit price. This points to the importance of not viewing shared-savings contracts as static, but rather of renegotiating the unit price when either party's costs change.

The key contribution of this paper lies in explicitly modeling demand *endogeneously* as a function of *effort* exerted *by both parties*, and analyzing how contract type can influence such endogeneous demand. As such, this paper opens up the largely unexplored field of contracting on indirect materials. Section 2 reviews relevant literature on supply-chain coordination, environmental management, and chemical supply contracts. Section 3 describes the basic model. Section 4 develops a general formulation that captures various types of contracts found in the literature. Section 5 derives equilibrium behavior under this general formulation. Section 6 discusses comparative statics, showing how these equilibria vary with the contract parameters. Section 7 describes the insights obtained from extensive numerical experimentation. Section 8 discusses practical implementation of shared-savings contracts and offers suggestions for further research. Section 9 summarizes our findings.

2. Literature Review

This paper draws on and contributes to several streams of literature reviewed below. It contributes to our understanding of the impact of contracts for indirect materials on supply-chain performance, using advanced concepts in game theory. It is also (to the best of our knowledge) the first attempt to rigorously model how contracts can reduce consumption, hence contributing to the environmental management literature.

A growing literature in supply-chain management, recently reviewed by Tsay et al. (1999), asks how redesigning contracts can lead to improved performance. Cachon and Zipkin (1999) show how cooperation between customer and supplier almost always leads to lower inventory costs than equilibrium behavior; we find analogous results here for total manufacturing and disposal costs. By focusing on indirect materials, we introduce an incentive conflict that is almost universal but fundamentally different from those studied so far: Suppliers earn more by selling more, while the customer earns more by using less.

We have heard various versions of the mantra that "the next step forward in environmental improvement lies in supply-chain coordination." A growing literature (see Thierry et al. 1995 and Fleischmann et al. 1997) focuses on the reverse flows induced by legislation requiring take-back of products and packaging. A different implication of supply-chain structure is the environmentally damaging incentive conflict mentioned earlier. Bierma and Waterstraat (1996, 2000) and Reiskin et al. (2000) describe a range of chemical management contracts used in practice to redress this; in this paper we analyze the implications of the different types of contract more precisely.

Because both parties can influence consumption and hence imperfectly internalize the benefits of their effort, this is reminiscent of the double moral hazard problem, discussed in Bhattacharyya and Lafontaine (1995) and Kim and Wang (1998). While that literature focuses only on profit optimization, we explicitly analyze the effects of specific contracts on both profits and consumption (or environmental impact).

Recent research has examined the link between environmental and economic performance. Klassen and McLaughlin (1996) find positive stock market effects after announcements of environmental awards and negative effects after environmental crises. This same argument is made more informally by Porter and van der Linde (1995) and Hart (1997). Our study contributes to this debate by formalizing a business model that is proving successful in practice, and which can also reduce environmental impacts.

3. The Basic Model

We first introduce the basic model and two benchmark cases, the volume-only base contract and the joint investment contract, which maximizes profits for the total chain. After that, we introduce and analyze shared-savings contracts in §§4 and beyond.

A single *supplier* produces (or purchases) some *indirect material* at a unit cost of c, and sells this to a single *customer*. The customer uses this indirect material in his own process, which generates revenue (net of all costs of manufacturing and raw materials other than the indirect material) of r per period, independently of the decisions modeled here. Also, assume for now that the unit price paid by the customer for the indirect material is $p > c$. The customer incurs costs d per unit due to material handling, disposal of waste material, etc. By exerting some effort, each party can reduce the amount of indirect material required per period. For example, either party might find ways to make the customer's production process more efficient, or the supplier might reformulate the product. Let e_s and e_c denote the amount of *use-reduction effort* exerted by the supplier and the customer, respectively, and let $0 \le e_s \le 1$ and $0 \le e_c \le 1$. Let $y(e_s, e_c)$ denote the quantity of indirect material required per period. Without loss of generality, we assume $y(0, 0) = 1$. Let $y(e_s, e_c)$ be twice continuously differentiable on $[0, 1] \times [0, 1]$. More effort leads to lower consumption, so that $(\partial y/\partial e_s)(e_s, e_c) < 0$ and $(\partial y/\partial e_c)(e_s, e_c) < 0$ for $(e_s, e_c) \in [0, 1] \times [0, 1]$. There are decreasing returns to effort, so that $(\partial^2 y/\partial e_s^2)(e_s, e_c) \ge 0$ and $(\partial^2 y/\partial e_c^2)(e_s, e_c) \ge 0$ for $(e_s, e_c) \in [0, 1] \times [0, 1]$. Occasionally, we will also require supermodularity of consumption in joint effort levels (which, in this case, is equivalent to nonnegative cross-derivatives: $(\partial^2 y/\partial e_c \partial e_s)(e_s, e_c) \ge 0$ for $(e_s, e_c) \in [0, 1] \times [0, 1]$).

Effort levels e_s and e_c cost the supplier and customer $c_s(e_s)$ and $c_c(e_c)$, respectively, per period; $c_s(e_s)$ and $c_c(e_c)$ are twice continuously differentiable on [0,1), and convex increasing in effort so that $c_s'(e_s) >$

0 and $c_c'(e_c) > 0$ for $(e_s, e_c) \in (0,1) \times (0,1)$, and $c_s''(e_s) > 0$ and $c_c''(e_c) > 0$ for $(e_s, e_c) \in [0,1) \times [0,1)$. We assume there is some "low-hanging fruit," i.e., that both parties initial marginal costs of effort are negligible, or $c_s'(0) = 0$ and $c_c'(0) = 0$. Finally, we assume $\lim_{e_s \to 1} c_s'(e_s) = \lim_{e_c \to 1} c_c'(e_c) = \infty$. The last two assumptions are not strictly necessary, but facilitate the exposition by ruling out border equilibria. Let Π_s and Π_c represent the profits to the supplier and customer, respectively, per period:

$$\Pi_s(e_s, e_c) = (p-c)y(e_s, e_c) - c_s(e_s), \tag{1}$$

$$\Pi_c(e_s, e_c) = r - (p+d)y(e_s, e_c) - c_c(e_c). \tag{2}$$

Also, let $\Pi_T \equiv \Pi_s + \Pi_c$ be the total channel profits. Initially, both players choose effort simultaneously; in our analysis of comparative statics we allow iterated play. In the *effort* game, the supplier selects e_s and the customer selects e_c (simultaneously) to maximize Π_s and Π_c, respectively. Let $e_s^*(e_c) = \arg\max_{e_s \in [0,1]} \Pi_s(e_s, e_c)$ and $e_c^*(e_s) = \arg\max_{e_c \in [0,1]} \Pi_c(e_s, e_c)$ be the best response functions for the supplier and customer, respectively. A pair $(\hat{e}_s, \hat{e}_c)$ is a Nash equilibrium if neither player can achieve higher profits by unilaterally changing his effort level, that is, $\Pi_s(\hat{e}_s, \hat{e}_c) \geq \Pi_s(e_s, \hat{e}_c)$ for all $e_s \in [0,1]$ and $\Pi_c(\hat{e}_s, \hat{e}_c) \geq \Pi_c(\hat{e}_s, e_c)$ for all $e_c \in [0,1]$.

Base Contract. The profit functions in (1) and (2) represent the traditional *base contract*. The supplier is compensated on a quantity basis: A fixed price per unit of indirect material sold. Clearly, the more he sells (at a given price), the higher his profit. The supplier has no incentive to exert any effort to reduce the quantity of indirect material needed. Since $\partial \Pi_s / \partial e_s \leq 0$ for all feasible (e_s, e_c), we have $e_s^*(e_c) = 0$ for all $e_c \in [0,1]$. The (componentwise) concavity assumptions on $y(e_s, e_c)$ and $c_c(e_c)$ establish concavity of $\Pi_c(e_s, e_c)$ in e_c, so the first-order condition is sufficient for characterizing the customer's best response function. In equilibrium $e_s = 0$ and the conditions on the cost-of-effort functions guarantee the existence of a unique $e_c \in (0,1)$ such that $\partial \Pi_c / \partial e_c = 0$, so that in the base contract, the unique effort equilibrium is $(0, \hat{e}_c^{BC})$, where $\hat{e}_c^{BC}$ is the unique e_c satisfying

$$c_c'(e_c) = -(p+d)\frac{\partial y}{\partial e_c}(0, e_c). \tag{3}$$

Although the supplier exerts no effort to reduce consumption, the customer will always exert some positive effort to this end. Each unit of indirect material causes the customer to incur a positive cost $p+d$, which is greater than the initial marginal cost of effort.

Joint Investment Contract. The base contract creates a situation where the supplier and customer act quite independently and with opposite incentives regarding consumption. Is it possible to better align the two players' incentives? To explore this, we consider the *joint investment* contract suggested by Bierma and Waterstraat (1996), essentially equivalent to a vertically integrated company.

Suppose the supplier and customer set up a semi-independent joint venture to which all costs related to the indirect material, including procurement costs, handling and disposal costs, and the cost of use-reduction efforts, are charged. The supplier pays a fraction λ and the customer pays a fraction $1-\lambda$ of the costs incurred by this joint venture; in addition, the customer pays the supplier some fixed transfer t. The costs incurred by the joint venture, and the profits earned by supplier and customer are, respectively,

$$C_{JV} = (c+d)y(e_s, e_c) + c_s(e_s) + c_c(e_c), \tag{4}$$

$$\Pi_s = t - \lambda C_{JV}, \quad \Pi_c = r - t - (1-\lambda)C_{JV}, \quad \Pi_T = r - C_{JV}. \tag{5}$$

The constants r and t do not affect the choice of effort levels. The resulting equilibrium effort levels represent the first-best solution, i.e., the effort levels that would be chosen by a centralized decision maker seeking to maximize total channel profits. The first-order optimality conditions for these optimal effort levels are given by

$$c_s'(e_s) = -(c+d)\frac{\partial y}{\partial e_s}(e_s, e_c) > 0, \tag{6}$$

$$c_c'(e_c) = -(c+d)\frac{\partial y}{\partial e_c}(e_s, e_c) > 0, \tag{7}$$

Compare Equations (6) and (7) to Equation (3). In the joint investment equilibrium characterized by (6), the supplier exerts some positive effort $\hat{e}_s^{JI}$. As a result of this, along with the fact that $p > c$ in the base contract, $\hat{e}_c^{JI} \leq \hat{e}_c^{BC}$. The base contract leads the supplier to underinvest and the customer to overinvest in effort, relative to the first-best solution. The joint investment contract requires that joint costs C_{JV} be shared, which also implies that both parties' costs of effort must be shared. Both players' effort levels and corresponding costs would typically be very difficult to verify, so this contract is usually impossible to implement in practice.

4. Shared-Savings Contracts

Since the base contract leads to suboptimal channel profits and joint investment is impractical, are there practical contracts that yield equilibria approaching or matching the first-best solution? Bierma and Waterstraat (1996) list several different contracts currently in use in the chemical sector. These arrangements, described below, can be applied to almost any indirect material. In §8, we revisit some schemes from practice and compare them to our modeling framework.

• *Chemical management fee.* The customer pays the supplier a fixed fee t that is independent of the volume of chemicals used. The supplier manages the customer's chemicals and takes responsibility for waste disposal. He passes chemical costs on to the customer at a nominal profit—say at some price p', with $p > p' > c$.

• *Leasing.* The customer leases the chemicals from the supplier at a unit price p'', for which he receives some management services from the supplier, who retains ownership of the chemicals and, with that, responsibility for waste management.

• *Shared savings.* The customer pays a fixed fee t, independent of the volume of chemicals used. Benefits from reduced chemical consumption (reductions in the supplier's chemical costs as well as the customer's handling and disposal costs) are shared, with the supplier receiving a fraction $\lambda \in [0, 1]$ and the customer the remaining fraction $(1-\lambda)$ of the total savings.

Each of the above arrangements can be modeled by replacing the original transfer py in our base contract with a two-part transfer $T(y)$. A chemical management fee corresponds to $T(y) = t + (p' - d)y$, a lease arrangement to $T(y) = (p'' - d)y$, and a shared-savings contract to $T(y) = t + \lambda(c+d)(1-y) - c(1-y)$. All the above contracts are of the form $T(y) = t + a \cdot y$. The parameter a redistributes the unit costs because of indirect material consumption, hopefully realigning incentives in a way that increases channel profits. The parameter t does not affect either player's effort decision directly, but can be adjusted to give each player the incentive to participate. Because any contract of this type can be seen as a mechanism for creating and then sharing cost savings, we use the term *shared savings* for all such contracts. We focus on the general form $T(y) = t + a \cdot y$, as this captures a broad class of contracts commonly found in practice, described in §8.

5. Equilibrium Efforts Under Shared-Savings Contracts

Given a transfer $T(y) = t + a \cdot y$, the players' and the joint profit functions are:

$$\Pi_s(e_s, e_c; a) = t + (a - c)y(e_s, e_c) - c_s(e_s), \tag{8}$$

$$\Pi_c(e_s, e_c; a) = r - t - (a + d)y(e_s, e_c) - c_c(e_c), \tag{9}$$

$$\Pi_T(e_s, e_c; a) = r - (c + d)y(e_s, e_c) - c_s(e_s) - c_c(e_c). \tag{10}$$

Because t does not influence equilibrium effort levels, we focus on the effect of a. If a is too large or too small, then one party will have no incentive to exert any effort to reduce consumption. Such extreme values of a always lead to suboptimal channel profits.

Proposition 1. *If $a \geq c$, then the unique equilibrium is $(0, \hat{e}_c^0(a))$, where $\hat{e}_c^0(a)$ is the unique e_c satisfying $c_c'(e_c) = -(a+d)(\partial y/\partial e_c)(0, e_c)$. In addition, $\hat{e}_c^0(a)$ is increasing in a and $\Pi_T(0, \hat{e}_c^0(a))$ is decreasing in a. If $a \leq -d$, then the unique equilibrium is $(\hat{e}_s^0(a), 0)$, where $\hat{e}_s^0(a)$ is the unique e_s satisfying $c_s'(e_s) = -(c-a)(\partial y/\partial e_s)(e_s, 0)$. In addition, $\hat{e}_s^0(a)$ is decreasing in a and $\Pi_T(\hat{e}_s^0(a), 0)$ is increasing in a.*

Proof. If $a \geq c$, then $\partial \Pi_s/\partial e_s = (a-c)(\partial y/\partial e_s)(e_s, e_c) - c_s'(e_s) \leq 0 \; \forall \; (e_s, e_c)$, so $e_s^*(e_c) = 0$ for all $e_c \in [0, 1]$. Arguments as for the base contract guarantee uniqueness of $\hat{e}_c^0(a)$. $\partial^2 \Pi_c/\partial e_c \partial a = -\partial y/\partial e_c > 0$, so Π_c is supermodular in (e_c, a) on the sublattice $[0, 1] \times [c, \infty)$, so that $\hat{e}_c^0(a)$ is increasing in a. The first-order condition defining $\hat{e}_c^0(a)$ gives $(d\Pi_T/da)(0, \hat{e}_c^0(a)) = (-(c + d)(\partial y/\partial e_c)(0, \hat{e}_c^0(a)) - c_c'(\hat{e}_c^0(a))) \times (d\hat{e}_c^0(a)/da)$

≤ 0. Similarly, if $a \leq -d$, then $\partial\Pi_c/\partial e_c \leq 0$, so $e_c^*(e_s) = 0$ for all $e_s \in [0, 1]$. Arguments as above guarantee uniqueness of $\hat{e}_s^0(a)$. $\hat{e}_s^0(a)$ is decreasing in a as $\partial^2\Pi_s/\partial e_s \partial a = \partial y/\partial e_s < 0$, so Π_s is submodular in (e_s, a) on the sublattice $[0, 1] \times (-\infty, -d]$. The first-order condition defining $\hat{e}_s^0(a)$ gives $(d\Pi_T/da)(\hat{e}_s^0(a), 0) \geq 0$. □

From the perspective of maximizing channel profits, we can restrict attention to $a \in [-d, c]$. If the supplier's profit margin $a - c$ is positive, he has no incentive to reduce consumption. Similar logic applies to the customer when $a < -d$. The base contract corresponds to $a = p > c$, so the above proposition implies that there always exists a shared savings contract with $-d \leq a \leq c$, for which channel profits will be at least as high as in the base contract.

If $-d \leq a \leq c$, the first-order optimality conditions for the supplier and customer are

$$c_s'(e_s) = -(c - a)\frac{\partial y}{\partial e_s}(e_s, e_c), \tag{10}$$

$$c_c'(e_c) = -(a + d)\frac{\partial y}{\partial e_c}(e_s, e_c). \tag{11}$$

The earlier assumptions on the cost functions guarantee that (10) and (11) have unique solutions, so that the best response functions, $e_s^*(e_c)$ and $e_c^*(e_s)$, are well defined. The following theorem establishes a number of fundamental facts about the best response functions and equilibrium effort levels, graphically illustrated in Figure 1.

Figure 1 **Best Response Functions with Multiple Equilibria**

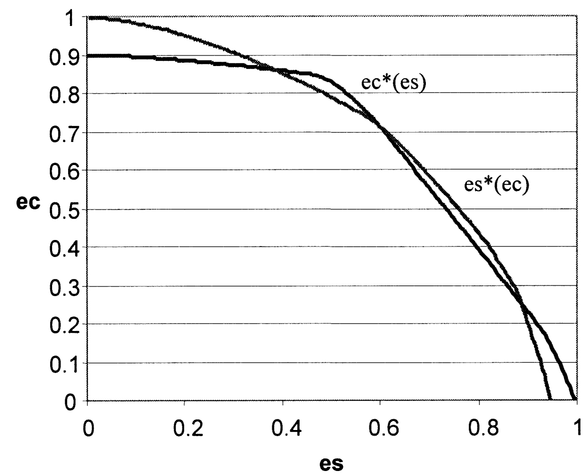

Note. The graph illustrates how, despite the smooth and well-behaved nature of the best response functions, one could obtain multiple equilibria.

THEOREM 1. *If $-d \leq a \leq c$, then the following are true*:

(i) *The best response functions $e_s^*(e_c)$ and $e_c^*(e_s)$ are both decreasing in their arguments.*

(ii) *There exists at least one equilibrium in the effort game.*

(iii) *The set of equilibria is a chain, i.e., if there are multiple equilibria, they can be ordered: for any equilibria $(\hat{e}_s, \hat{e}_c)$ and $(\bar{e}_s, \bar{e}_c)$, either $\hat{e}_s \leq \bar{e}_s$ and $\hat{e}_c \geq \bar{e}_c$, or $\hat{e}_s \geq \bar{e}_s$ and $\hat{e}_c \leq \bar{e}_c$.*

(iv) *If there are multiple equilibria, then for any equilibria $(\hat{e}_s, \hat{e}_c)$ and $(\bar{e}_s, \bar{e}_c)$ ordered so that $\hat{e}_s \leq \bar{e}_s$ and $\hat{e}_c \geq \bar{e}_c$, we have $\Pi_s(\hat{e}_s, \hat{e}_c) \geq \Pi_s(\bar{e}_s, \bar{e}_c)$ and $\Pi_c(\hat{e}_s, \hat{e}_c) \leq \Pi_c(\bar{e}_s, \bar{e}_c)$.*

PROOF. For (i) and (ii), use the change of variable $\tilde{e}_s = -e_s$. Then, $\partial^2\Pi_s/\partial\tilde{e}_s\partial e_c = (c - a)(\partial^2 y/\partial e_s \partial e_c) \geq 0$ and $\partial^2\Pi_c/\partial\tilde{e}_s\partial e_c = (a + d)(\partial^2 y/\partial e_s \partial e_c) \geq 0$, so Π_s and Π_c are both supermodular in $(\tilde{e}_e, e_c)$ on the lattice $[-1, 0] \times [0, 1]$ if and only if y is supermodular, which we assumed earlier. Theorem 1.2 in Topkis (1979) immediately establishes (i), and Theorem 3.1 in Topkis (1979) establishes (ii). For (iii), consider two equilibria $(\hat{e}_s, \hat{e}_c)$, and $(\bar{e}_s, \bar{e}_c)$, assume without loss of generality that $\hat{e}_s \leq \bar{e}_s$. Since each player's best response function is decreasing in the other player's effort, $\hat{e}_c = e_c^*(\hat{e}_s) \geq e_c^*(\bar{e}_s) = \bar{e}_c$. Finally, for Part (iv), note that $\Pi_s(\bar{e}_s, \bar{e}_c) \leq \Pi_s(\bar{e}_s, \hat{e}_c) \leq \Pi_s(e_s^*(\hat{e}_c), \hat{e}_c) = \Pi_s(\hat{e}_s, \hat{e}_c)$ and $\Pi_c(\hat{e}_s, \hat{e}_c) \leq \Pi_c(\bar{e}_s, \hat{e}_c) \leq \Pi_c(\bar{e}_s, e_c^*(\bar{e}_s)) = \Pi_c(\bar{e}_s, \bar{e}_c)$. □

Result (i) states that the use-reduction efforts exerted by the players act as substitutes—the more effort one player exerts, the less the other one will. Comparing any two equilibria, each player earns higher profits at the equilibrium in which he exerts lower effort. Recall that our goal is to achieve equilibrium effort levels that yield higher channel profits than the base contract and, if possible, match the joint investment effort levels. Although Proposition 1 and Theorem 1 show that the former goal can be achieved, shared savings contracts cannot achieve full channel coordination. To see this, compare the first-order conditions (10) and (11) for the shared savings contract with (6) and (7) for joint investment. Let $e_s^{SS}(e_c)$, $e_c^{SS}(e_s)$, $e_s^{JI}(e_c)$, $e_c^{JI}(e_s)$ be the best response functions under shared savings and under joint investment, respectively. Since $-d \leq a \leq c$, (6) and (10) imply that $e_s^{SS}(e_c) \leq e_s^{II}(e_c)$, while (7) and (11) imply that $e_c^{SS}(e_s) \leq e_c^{II}(e_s)$, i.e., faced with any given effort level by his counterpart, each player will exert less effort under a shared-savings contract than under vertical integration. In addition, since $a < c$ implies $e_c^{SS}(e_s) < e_c^{II}(e_s)$ and $a > -d$ implies $e_s^{SS}(e_c) < e_s^{II}(e_c)$, no shared-savings contract can induce the first-best solution. The numerical experiments in §7 do suggest, however, that simple shared-savings contracts can lead to substantial improvements over the base contract.

6. Comparative Statics of Equilibria

Above we established existence of equilibria in the effort game, for given parameters a, c, and d. We now explore how the equilibria behave as these parameters change; we first study comparative statics of effort levels, then of consumption, and finally of total profits.

6.1. Comparative Statics of Effort Levels in Equilibrium

Comparative statics is challenging as there may be multiple equilibria. We use the concept of fictitious or iterated play from Lippman et al. (1987). (Note that iterated play is quite different from iterated dominance; see Fudenberg and Tirole 1991). Let the supplier pick any effort level e_s in [0,1]. The customer chooses his best response $e_c^*(e_s)$, to which the supplier's best response is $e_s^*(e_c^*(e_s))$. Constructing an infinitely iterated sequence, play always converges to an equilibrium. Some equilibria are "attractors," in that iterated play starting from either side sufficiently close to that equilibrium will converge to it; other equilibria are "repellent," in that iterated play will always diverge away. Assume that for some given contract a_1, the players have converged to an arbitrary equilibrium $e^1(a_1) = (e_s^1(a_1), e_c^1(a_1))$. Now change the contract, such that $a_2 > a_1$. The former equilibrium e^1 is no longer on either party's best-response curve, so iterated play starts anew, converging to a new attracting equilibrium. Loosely speaking, Proposition 2 states that as a increases, all attracting equilibria shift towards lower supplier and higher customer effort. Proofs from here onwards are provided in the Appendix.

Proposition 2. *For any equilibrium $e^1(a_1)$, pick $a_2 > a_1$. Iterated play starting from $e^1(a_1)$ but under contract a_2 will converge to a new equilibrium, $e^2(a_2)$, such that $e_s^2(a_2) \leq e_s^1(a_1)$ and $e_c^2(a_2) \geq e_c^1(a_1)$. Similarly, if $a_2 < a_1$, then $e_s^2(a_2) \geq e_s^1(a_1)$ and $e_c^2(a_2) \leq e_c^1(a_1)$.*

Analogous results for changes in c and d are stated in Proposition 3.

Proposition 3. *For any equilibrium $e^1(c_1)$, pick $c_2 > c_1$. Iterated play starting from $e^1(c_1)$ but under cost c_2 will converge to a new equilibrium, $e^2(c_2)$, such that $e_s^2(c_2) \geq e_s^1(c_1)$ and $e_c^2(c_2) \leq e_c^1(c_1)$. Similarly, if $c_2 < c_1$, then $e_s^2(c_2) \leq e_s^1(c_1)$ and $e_c^2(c_2) \geq e_c^1(c_1)$. For any equilibrium, $e^1(d_1)$, pick $d_2 > d_1$. Iterated play starting from $e^1(d_1)$ but under cost d_2 will converge to a new equilibrium, $e^2(d_2)$, such that $e_s^2(d_2) \leq e_s^1(d_1)$ and $e_c^2(d_2) \geq e_c^1(d_1)$. Similarly, if $d_2 < d_1$, then $e_s^2(d_2) \geq e_s^1(d_1)$ and $e_c^2(d_2) \leq e_c^1(d_1)$.*

These results are intuitive, but it is less obvious that a change in a, c, or d will *always* have the expected effect, even if there are multiple equilibria. Indeed, the comparative statics for consumption and profits do not always behave as one would expect.

6.2. Comparative Statics of Consumption Levels in Equilibrium

We have established that an increase in one party's cost of consumption will provoke a shift of effort to that party. How do such cost changes affect consumption and total profits? One would expect that higher unit consumption costs should lead to reduced consumption after adjustment to the new equilibrium effort levels. We have not been able to guarantee that this will hold, as we are not dealing with (jointly) optimal effort levels but with equilibria in the effort game. Fortunately, one can always adjust a such that changes in c or d do have the expected effects. To show this we must assume uniqueness of the equilibrium in the effort game. Some sufficient conditions for uniqueness do exist but they are quite restrictive; see for instance Vives (1999, p. 47). Uniqueness was established theoretically for some of our numerical experiments and was confirmed numerically for the others. Let $\hat{e}(c, d, a) \equiv$

$(\hat{e}_s(c, d, a), \hat{e}_c(c, d, a))$ represent this unique equilibrium for any given vector (c, d, a). Consider a particular benchmark parameter vector (c_1, d_1, a_1), its associated equilibrium $\hat{e}(c_1, d_1, a_1)$, and the resulting consumption $y(\hat{e}_s(c_1, d_1, a_1), \hat{e}_c(c_1, d_1, a_1))$.

Proposition 4. *For any $c_2 > c_1$, there exists $a_2 \in [a_1, a_1 + c_2 - c_1]$ such that, at the new equilibrium $\hat{e}(c_2, d_1, a_2)$, consumption is lower than before, i.e., $y(\hat{e}_s(c_2, d_1, a_2), \hat{e}_c(c_2, d_1, a_2)) \leq y(\hat{e}_s(c_1, d_1, a_1), \hat{e}_c(c_1, d_1, a_1))$. Similarly, for any $d_2 > d_1$, there exists $a_2 \in [a_1 - (d_2 - d_1), a_1]$ such that, at the new equilibrium $\hat{e}(c_1, d_2, a_2)$, consumption is lower than before, i.e., $y(\hat{e}_s(c_1, d_2, a_2), \hat{e}_c(c_1, d_2, a_2)) \leq y(\hat{e}_s(c_1, d_1, a_1), \hat{e}_c(c_1, d_1, a_1))$.*

This result not only establishes the existence of an adjusted contract that yields the desired results, it also identifies the direction of adjustment needed. An increase in c can be adjusted for by increasing a, while an increase in d calls for a reduction in a. These adjustments help mitigate the cost increase by sharing it between both players, consistent with the spirit of shared-savings contracts.

6.3. Comparative Statics of Total Profits in Equilibrium

Although an increase in costs could lead to higher consumption, one can always adjust the contract such that consumption will decrease. If consumption can be written as $y(e_s, e_c) = y_s(e_s)y_c(e_c)$, we can show that equivalent results hold for total profits. This multiplicative form arises when the effects of both parties' efforts are independent. (This condition may not be necessary, but we have not been able to show the result without it.) As before, consider a benchmark vector (c_1, d_1, a_1) and its associated equilibrium $\hat{e}(c_1, d_1, a_1)$. For any given (c, d, a), let $\Pi_s(c, d, a)$, $\Pi_c(c, d, a)$, and $\Pi_T(c, d, a)$ represent the supplier's, customer's, and channel profits, respectively, which occur at the corresponding unique equilibrium.

Proposition 5. *For any $c_2 < c_1$, there exists $a_2 \in [a_1 - (c_1 - c_2), a_1]$ such that, at the new equilibrium $\hat{e}(c_2, d_1, a_2)$, channel profits are higher than before the parameter changes, i.e., $\Pi_T(c_2, d_1, a_2) \geq \Pi_T(c_1, d_1, a_1)$. Similarly, for any $d_2 < d_1$, there exists $a_2 \in [a_1, a_1 + (d_1 - d_2)]$ such that, at the new equilibrium $\hat{e}(c_1, d_2, a_2)$, channel profits are higher than before the parameter changes, i.e., $\Pi_T(c_1, d_2, a_2) \geq \Pi_T(c_1, d_1, a_1)$.*

7. Numerical Examples

In this section, we describe a numerical study that examines the profit and consumption performance of shared-savings contracts relative to the base contract and to the joint investment case. We use two functional forms for the consumption function (one additive and one multiplicative), each with a number of parameter combinations: $y_1(e_s, e_c) = (1 - e_s)^{\alpha_s}(1 - e_c)^{\alpha_c}$, with $\alpha_s, \alpha_c \in \{1, 2\}$, and $y_2(e_s, e_c) = (\alpha_s + \alpha_c - \alpha_s e_s - \alpha_c e_c)/(\alpha_s + \alpha_c)$, with $\alpha_c = 1$, $\alpha_s \in \{0.1, 0.2, 1, 5, 10\}$. Since $y_2(e_s, e_c)$ depends only on the ratio of the two parameters, α_c is fixed. A sufficient condition for the best-response map to be a contraction, and hence to ensure uniqueness of the equilibrium, is given in Vives (1999, p. 47):

$$\frac{\partial^2 \Pi_c}{\partial e_c^2} + \left|\frac{\partial^2 \Pi_c}{\partial e_c \partial e_s}\right| < 0 \quad \text{and} \quad \frac{\partial^2 \Pi_s}{\partial e_s^2} + \left|\frac{\partial^2 \Pi_s}{\partial e_c \partial e_s}\right| < 0 \quad \text{for all } e_s, e_c \in [0, 1].$$

For $y_2(e_s, e_c)$, and for $y_1(e_s, e_c)$ with $\alpha_s = \alpha_c = 1$, this condition is always met.

We also use two functional forms for the cost-of-effort functions: $c_1(e_s) = (1 - e_s)^{-x_s} - x_s e_s - 1$, $c_1(e_c) = (1 - e_c)^{-x_c} - x_c e_c - 1$, with $x_s, x_c \in \{1, 3\}$, and $c_2(e_s) = x_s(-\ln(1 - e_s) - e_s)$, $c_2(e_c) = x_c(-\ln(1 - e_c) - e_c)$, with $x_s, x_c \in \{10, 50\}$. (The first tends to be more "L-shaped," while the second is more "U-shaped.") In any given scenario, the supplier's and customer's cost-of-effort functions are of the same form, but all four combinations of parameter values are used. In all scenarios considered, we fix the revenue parameter at $r = 50$; other parameter values are $c \in \{5, 10\}$, $d \in \{0.25c, c, 4c\}$, and $p \in \{1.2c, 2c\}$.

We first present two specific examples from the study to illustrate counterintuitive behavior that can occur in certain settings.

Example 1. Proposition 5 stated that, faced with a decrease in c, it is always possible to adjust the contract so that channel profits increase. Channel profits could actually decrease without such an adjustment. Consider $c_1(e_s)$, $c_1(e_c)$, and $y_1(e_s, e_c)$ (defined above), and parameters $x_s = 1, x_c = 3, \alpha_s = \alpha_c = 1, r = 50, c = 10, d = 2.5$, and $a = 8$. The unique equilibrium is $(\hat{e}_s, \hat{e}_c) = (0.367, 0.253)$, and channel profits are $\Pi_T = 43.23$. If c decreases to $c = 8.5$ and the contract is not adjusted, the new equilibrium is $(\hat{e}_s, \hat{e}_c) = (0.140, 0.293)$, and channel profits drop to $\Pi_T = 42.34$. However, if a is reduced from 8 to 7, sharing some (but not all) of the supplier's cost decrease with the customer, the new equilibrium is $(\hat{e}_s, \hat{e}_c) = (0.314, 0.251)$, and channel profits are $\Pi_T = 43.58$. (Note that $a = 8$ is not the optimal contract for the original setting. While we were unable to prove that this behavior cannot occur when starting at the optimal contract, we were also unable to identify examples where it does occur. This suggests that the system may be more "well behaved" near the optimal contract.)

Example 2. Although proponents of shared-savings contracts argue that such arrangements lead to higher profits and reduced consumption (relative to the base contract), this need not be the case. Consider the cost-of-effort and consumption functions as in Example 1 and parameters $x_s = 3, x_c = 1, \alpha_s = \alpha_c = 1, r = 50, c = 10$, and $d = 2.5$. To allow comparison to a base contract with $p = 20$, we consider values for a outside the range $-d \le a \le c$. Figure 2 shows that channel profits are maximized at $a = 6.5$. When $a \le 6.5$, consumption decreases as channel profits increase; unfortunately, when $a \ge 6.5$, higher channel profits only come at the expense of higher consumption. Even more disappointing: Consumption is minimized at $a = 20$, corresponding to the base contract. So, although a shared-savings contract increases channel profits, it also increases consumption.

Figure 2 **Channel Profits and Consumption As a Function of *a***

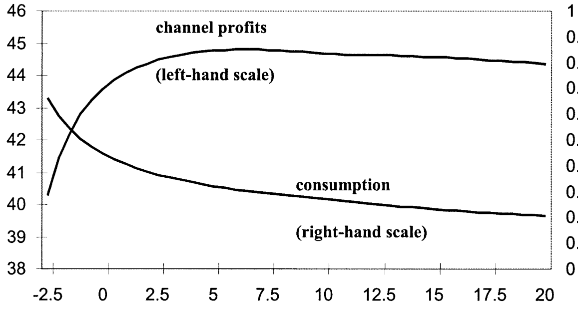

For the complete numerical study, we combine all functional forms and parameter values to yield 864 scenarios. For each scenario we vary a throughout the range $-d \le a \le c$, calculate the equilibrium effort levels, and select the contract that yields the highest channel profit at equilibrium, denoted by Π_{SS}. We also compute the equilibrium effort levels under the base contract and under joint investment (i.e., the first-best effort levels); denote the corresponding channel profits by Π_{BC} and Π_{FB}, respectively. We report the extent to which the selected shared-savings contract closes the profit gap between the base contract and the first-best case:

$$\text{shared-savings capture percentage (SSCP)} = \left(\frac{\Pi_{SS} - \Pi_{BC}}{\Pi_{FB} - \Pi_{BC}}\right) \times 100\%.$$

The numerical study led to the following tentative observations. First, a unique equilibrium exists for the optimal shared-savings contract in every one of the 864 scenarios considered, even when the sufficient condition above was not met. (Uniqueness can be verified by iterated play starting at $(e_s, e_c) = (0, 1)$ and at $(e_s, e_c) = (1, 0)$. If both sequences converge to the same equilibrium, that equilibrium is unique; see Topkis 1979, Theorem 4.2.) The apparent prevalence of a unique equilibrium helps make shared-savings contracts practical to implement. It also lends support to the uniqueness assumption made in §§6.2 and 6.3.

Second, the optimal shared-savings contract tends to capture a significant percentage of the maximum potential profit gain. The unweighted average of SSCP across all 864 scenarios is 73.07%. Shared-savings contracts seem to perform better in percentage terms when the maximum potential gain is large. The aggregate shared-savings capture percentage (the sum of the shared-savings gains across all scenarios divided by the sum of the maximum potential gains across all scenarios) is even higher, at 90.82%. Scenarios where the supplier is relatively more effective in reducing consumption (high α_s or low α_c values) tend to yield higher SSCPs, since supplier effort induced by the shared-savings contract causes a greater consumption reduction.

Third, the effects on consumption vary significantly across scenarios, ranging from a decrease of 89% to an increase of 82%. The unweighted average of the shared-savings consumption-reduction percentages (versus the base contract) across all 864 scenarios is 9.05%. Again, shared-savings contracts seem to perform better in percentage terms when the potential gain is large. The aggregate shared-savings consumption-reduction percentage (the sum of the shared-savings consumption reductions across all scenarios divided by the sum of the base contract consumption levels across all scenarios) is higher, at 18.20%. Again, scenarios where the supplier is relatively more effective in reducing consumption (high α_s or low α_c values) tend to yield greater consumption-reduction percentages.

Finally, the correlation between shared-savings capture percentages and consumption-reduction percentages is fairly high at 0.63, so that superior profit performance tends to be associated with positive consumption reductions. Cases with a high p, by contrast, tend to result in a high shared-savings capture percentage, but a low or even negative consumption reduction, as customer effort decreases without a compensating increase in supplier effort. Table 1 shows results for 24 representative scenarios. In the second set of results (with $\alpha_s = 2$), the supplier's effort has more impact than in the first set (with $\alpha_s = 1$); the results illustrate that the value of shared savings is greater with respect to profits, but especially with respect to consumption reduction. The table also illustrates the other observations made above.

8. Discussion: Shared-Savings Contracts in Practice and Areas for Future Research

The theoretical and numerical analyses in this paper suggest that shared-savings contracts have the potential to improve coordination within a supply chain (even though full coordination is ruled out), but one may wonder: Are such contracts feasible in practice? The answer is a resounding yes. Shared-savings contracts were recently featured on National Public Radio (April 26, 2000) and advertised in *Fortune* magazine (May 29, 2000), and the Chemical Strategies Partnership offers extensive information on other applications. In this section, we discuss some arrangements currently in use and how they relate to the shared-savings contracts studied here. The first set of examples are discussed in Bierma and Waterstraat (2000); the next two were selected from a series of interviews and site visits by the authors.

The Total Fluids Management and Total Solvents Management programs between the Ford Chicago Assembly Plant and PPG's Chemfil division specify a fixed fee per vehicle based on historical chemical usage, and a fixed annual fee for chemicals that cannot be linked to production volume. These arrangements correspond to contracts with positive t and $a \in [-d, 0]$ in our model, depending on the degree to which PPG manages and pays for Ford's handling and disposal costs. PPG's Pay-as-Painted program with Chrysler's Belvedere Assembly Plant in Belvedere, Illinois, is highly similar. Two arrangements combining fixed fees per vehicle with management fees for selected services are those between GM's Truck and Bus Plant in Janesville, Wisconsin, and BetzDearborn, and between GM's Electro-Motive Division in LaGrange, Illinois, and D.A. Stuart Company.

Another chemicals supplier offers a three-part arrangement: A flat service fee, reimbursement for chemicals used, and gainsharing on commodities and chemical efficiency improvement. This adds a reimbursement component cy to the "shared-savings" contract mentioned in §4: $T(y) = t + cy + \lambda(c+d)(1-y) - c(1-y) = t + cy + ((1-\lambda)c - \lambda d)y - ((1-\lambda)c - \lambda d)$.

A US-based chemical manufacturer transformed itself during the 1990s, changing from selling exclusively its own products, on a volume basis, to acting primarily as a service provider whose own products only account for 5% of sales and who thrives on fixed cost programs, i.e., $T(y) = t + ay$ as in several other examples above.

The model proposed here is a starting point for studying a wide range of contracts for procurement of indirect materials. It captures some key incentive conflicts not studied before in the literature, while

Table 1 Associations Among Potential Profit Gain, Shared-Savings Capture Percentage, and Consumption Reduction Percentage

	α_s	c	d	p	Max Potential Profit Gain $\Pi_{FB} - \Pi_{BC}$	Shared-Savings Capture Percentage (%)	Consumption Reduction Percentage (%)
Scenario 1	1	5	1.25	6	0.3219	48.61	−12.02
Scenario 2	1	5	5	6	0.5493	53.18	−9.74
Scenario 3	1	5	20	6	1.4758	63.43	−6.24
Scenario 4	1	5	1.25	10	0.5066	67.34	−29.70
Scenario 5	1	5	5	10	0.6720	61.73	−22.26
Scenario 6	1	5	20	10	1.5224	64.55	−12.04
Scenario 7	1	10	2.5	12	0.7185	56.46	−11.38
Scenario 8	1	10	10	12	1.1775	61.16	−9.37
Scenario 9	1	10	40	12	2.9518	70.45	−2.66
Scenario 10	1	10	2.5	20	1.0303	69.64	−30.44
Scenario 11	1	10	10	20	1.3798	66.86	−22.63
Scenario 12	1	10	40	20	3.0259	71.17	−8.44
Average					1.2777	62.88	−14.74
	α_s	c	d	p	Max Potential Profit Gain $\Pi_{FB} - \Pi_{BC}$	Shared-Savings Capture Percentage (%)	Consumption Reduction Percentage (%)
Scenario 1	2	5	1.25	6	0.8884	76.99	3.76
Scenario 2	2	5	5	6	1.4599	80.00	8.53
Scenario 3	2	5	20	6	3.5707	85.12	17.87
Scenario 4	2	5	1.25	10	1.0730	80.95	−11.42
Scenario 5	2	5	5	10	1.5826	81.55	−1.91
Scenario 6	2	5	20	10	3.6173	85.31	13.40
Scenario 7	2	10	2.5	12	1.8455	81.52	8.54
Scenario 8	2	10	10	12	2.9044	84.04	14.16
Scenario 9	2	10	40	12	6.6471	88.12	24.34
Scenario 10	2	10	2.5	20	2.1573	84.19	−7.11
Scenario 11	2	10	10	20	3.1067	85.08	3.75
Scenario 12	2	10	40	20	6.7211	88.25	20.08
Average					2.9645	83.43	7.83

Note. All scenarios use $y_1(e_s, e_c) = (1-e_s)^{\alpha_s}(1-e_c)^{\alpha_c}$, $c_1(e_s) = (1-e_s)^{-x_s} - x_s e_s - 1$, $c_1(e_c) = (1-e_c)^{-x_c} - x_c e_c - 1$, $\alpha_c = 2$, $x_s = x_c = 3$, and $r = 50$.

obviously being a simplification. Contracts in practice display significant variety and complexity, though often with the shared-savings contract studied here as a fundamental ingredient. Some features of these contracts which we have not captured here do pose interesting questions for future research.

First, a key benefit of chemical management services programs to suppliers lies in ensuring continuing business with its customers. In such instances, the supplier exerts effort now, expecting more profits in future. Second, in practice, two types of cost reduction occur: consumption reduction (e.g., by reducing waste or by improving the customers' process) and unit cost reduction (e.g., through order consolidation or substituting less expensive chemicals for more costly ones). The early cost savings are often achieved by the supplier's efforts alone and are primarily unit cost reductions, whereas consumption reductions (which require both supplier and customer effort) follow when the relationship has become more mature. A more sophisticated cost and consumption function would be needed to capture these effects.

Third, we have presented the most general framework possible here; it would be valuable to identify specific consumption and cost-of-effort functions to obtain more specific guidelines on how to set the contract parameter a. Another extension would be to make the dependence of consumption on effort stochastic. Current work is ongoing in both directions, building on the double moral hazard literature in economics (see, e.g., Kim and Wang 1998).

9. Summary and Conclusions

In this paper, we have studied how contracts affect consumption of indirect materials, by influencing the amount of effort supplier and customer exert to reduce that consumption. We compared a variety of contracts currently found in industry, including the base contract in which the supplier earns a positive margin on each unit sold, joint investment in which both parties choose the jointly optimal effort levels, and shared-savings contracts, in which both parties would benefit from any reduction in consumption. We have shown that, contrary to common belief, such shared-savings contracts can always lead to higher channel profits, but not necessarily to lower consumption. In our numerical experiments appropriate shared-savings contracts achieve, on average, 73% of the potential profit improvement. We have also shown that modifying the unit price specified in the contract will unambiguously cause one party's effort to be substituted by the other, even though there may be multiple effort equilibria. Finally, and again unexpectedly, a decrease in either party's cost could lead to a decrease in channel profits; this is because of the relatively "unpredictable" nature of studying comparative statics of equilibria, rather than sensitivity analysis of an optimum. When costs change, it might be necessary but is always possible to adjust the contract in order for consumption and profits to behave as expected. This paper has, for the first time, rigorously established the value of shared-savings contracts for indirect materials, while simultaneously pointing to some as yet unrecognized potential pitfalls in their implementation.

Acknowledgments

The first author received financial support from the AT&T Foundation for Industrial Ecology. The authors are also grateful to the referees and associate editor for their helpful comments on an earlier version, to Professor Bierma for sharing his thoughts with us, to all interviewees at the chemical management service providers and other firms for their openness, and to Pierre-Laurent Wetli for his research assistance during the interviews.

Appendix. Proofs of Comparative Statics Results in Propositions 2–5

Proof of Proposition 2. The proof relies largely on Theorems A and B in Lippman et al. (1987), for which both best-response functions must be decreasing; Theorem 1 already showed this. The composite best response function $f_s(\tilde{e}_s; a)$, defined by $f_s(\tilde{e}_s; a) = \tilde{e}_s^*(e_c^*(\tilde{e}_s; a); a)$, must be isotone in $\tilde{e}_s$, where $\tilde{e}_s := -e_s$ and $\tilde{e}_s^*(e_c) := -e_s^*(e_c)$. (Isotonicity is the generalization of monotonically increasing to functions on lattices.) This follows from the fact that the function $\tilde{e}_s^*(e_c)$ is increasing in e_c. Tarski's (1955) fixed-point theorem shows that any isotone function from a complete lattice to itself has a fixed point; this ensures existence of equilibria. Finally, $f_s(\tilde{e}_s; a)$ must be isotone in a, which holds as $e_c^*(\tilde{e}_s; a)$ is increasing in a and $\tilde{e}_s^*(e_c; a)$ is increasing in e_c and in a. Let $F(e; a)$ denote the fixed point towards which iterated play starting at e converges when the contract is a. Now, for any equilibrium $\tilde{e}^1(a_1) = (\tilde{e}_s^1(a_1), e_c^1(a_1))$, we must have $F(\tilde{e}^1(a_1); a_2) \geq \tilde{e}^1(a_1)$, which is precisely what we set out to show. The arguments can easily be duplicated for the case $a_2 < a_1$. □

Proof of Proposition 3. Since $\partial^2\Pi_s/\partial e_s \partial c = -\partial y/\partial e_s > 0$ and $\partial^2\Pi_c/\partial e_c \partial c = 0$, it follows that $e_s^*(\tilde{e}_c; c)$ is increasing in $\tilde{e}_c$ and in c, and $\tilde{e}_c^*(e_s; c)$ is increasing in e_s, where $\tilde{e}_c := -e_c$ and $\tilde{e}_c^*(e_s) := -e_c^*(e_s)$. Similarly, $\partial^2\Pi_s/\partial e_s \partial d = 0$ and $\partial^2\Pi_c/\partial e_c \partial d = -(\partial y/\partial e_c) > 0$, so $\tilde{e}_s^*(e_c; d)$ is increasing in e_c, and $e_c^*(\tilde{e}_s; c)$ is increasing in $\tilde{e}_s$ and in d. The rest of the proof mirrors that of Proposition 2 and is omitted. □

Proof of Proposition 4. Consider the equilibrium $\hat{e}_s(c_2, d_1, a_1)$ resulting from $c_2 > c_1$ and the original contract a_1. Proposition 3 implies that $\hat{e}_s(c_2, d_1, a_1) \geq \hat{e}_s(c_1, d_1, a_1)$ and $\hat{e}_c(c_2, d_1, a_1) \leq \hat{e}_c(c_1, d_1, a_1)$. We seek a_2 such that $\hat{e}_c(c_2, d_1, a_2) = \hat{e}_c(c_1, d_1, a_1)$. Note that Proposition 2 and the inequalities above imply that $a_2 \geq a_1$. Now consider $a = a_1 + c_2 - c_1$. Under this contract the supplier's and customer's first-order optimality conditions become:

$$c_s'(e_s) = (a - c_2)\frac{\partial y(e_s, e_c)}{\partial e_s} = ((a_1 + c_2 - c_1) - c_2)\frac{\partial y(e_s, e_c)}{\partial e_s} = (a_1 - c_1)\frac{\partial y(e_s, e_c)}{\partial e_s}$$

$$c_c'(e_c) = -(a + d_1)\frac{\partial y(e_s, e_c)}{\partial e_c} = -(a_1 + c_2 - c_1 + d_1)\frac{\partial y(e_s, e_c)}{\partial e_c}.$$

Since these conditions are identical to those under (c_1, d_2, a_1), where $d_2 = d_1 + c_2 - c_1 > d_1$, Proposition 3 implies $\hat{e}_s(c_2, d_1, a_1 + c_2 - c_1) \leq \hat{e}_s(c_1, d_1, a_1)$ and $\hat{e}_c(c_2, d_1, a_1 + c_2 - c_1) \geq \hat{e}_c(c_1, d_1, a_1)$. Therefore one such a_2 satisfies $a_1 \leq a_2 \leq a_1 + c_2 - c_1$. We now wish to show that $\hat{e}_s(c_2, d_1, a_2) \geq \hat{e}_s(c_1, d_1, a_1)$. To that end, note that the supplier's profit function can be written as $\Pi_s = t + ky(e_s, e_c) - c_s(e_s)$, where

$$k = \begin{cases} a_1 - c_1, & \text{for parameters } (c_1, d_1, a_1) \\ a_2 - c_2 \leq a_1 - c_1, & \text{for parameters } (c_2, d_1, a_2) \end{cases}.$$

Since $\partial^2\Pi_s/\partial e_s \partial k = \partial y/\partial e_s < 0$, $e_s^*(e_c; k)$ is decreasing in k. Therefore, $\hat{e}_s(c_2, d_1, a_2) = e_s^*(\hat{e}_c(c_2, d_1, a_2);\ \ a_2 - c_2) = e_s^*(\hat{e}_c(c_1, d_1, a_1);\ a_2 - c_2) \geq e_s^*(\hat{e}_c(c_1, d_1, a_1);\ a_1 - c_1) = \hat{e}_s(c_1, d_1, a_1)$. So $y(\hat{e}_s(c_2, d_1, a_2), \hat{e}_c(c_2, d_1, a_2)) \leq y(\hat{e}_s(c_1, d_1, a_1), \hat{e}_c(c_1, d_1, a_1))$. $\square$

The proof of Proposition 5 is similar in flavor, and is available from the authors, or on the *Management Science* website ⟨http://mansci.pubs.informs.org⟩.

References

Bhattacharyya, S., F. Lafontaine. 1995. Double-sided moral hazard and the nature of share contracts. *RAND J. Econom.* **26** 761–781.

Bierma, T.J., F.L. Waterstraat Jr. 1996. P2 assistance from your supplier. *Pollution Prevention Rev.* **Autumn** 13–24.

——, ——. 2000. *Chemical Management: Reducing Waste and Cost Through Innovative Supply Strategies*. John Wiley and Sons, New York.

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

Fleischmann, M., J.M. Bloemhof-Ruwaard, R. Dekker, E. van der Laan, J.A.E.E. van Nunen, L.N. Van Wassenhove. 1997. Quantitative models for reverse logistics: A review. *Eur. J. Oper. Res.* **103** 1–17.

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

Hart, S.L. 1997. Beyond greening: Strategies for a sustainable world. *Harvard Bus. Rev.* **75**(1) 66–76.

Kim, S.K., S. Wang. 1998. Linear contracts and the double moral-hazard. *J. Econom. Theory* **82** 342–378.

Klassen, R.D., C. McLaughlin. 1996. The impact of environmental management on firm performance. *Management Sci.* **42**(8) 1199–1214.

Lippman, S.A., J.W. Mamer, K.F. McCardle. 1987. Comparative statics in non-cooperative games via transfinitely iterated play. *J. Econom. Theory* **41** 288–307.

Porter, M.E., C. van der Linde. 1995. Green and competitive: Ending the stalemate. *Harvard Bus. Rev.* **73**(5) 120–134.

Reiskin, E.D., A.L. White, J.K. Johnson, T.J. Votta. 2000. Servicizing the chemical supply chain. *J. Indust. Ecology* **3**(2&3) 19–31.

Tarski, A. 1955. A lattice-theoretical fixpoint theorem and its applications. *Pacific J. Math.* **5** 285–309.

Thierry, M., M. Salomon, J. van Nunen, L. Van Wassenhove. 1995. Strategic issues in product recovery management. *California Management Rev.* **37**(2) 114–135.

Topkis, D.M. 1979. Equilibrium points in nonzero-sum *n*-person submodular games. *SIAM J. Control Optim.* **17**(6) 773–787.

Tsay, A.A., S. Nahmias, N. Agrawal. 1999. Modeling supply chain contracts: A review. S. Tayur, R. Ganeshan, M. Magazine, eds. *Quantitative Models for Supply Chain Management.* Kluwer Academic Publishers, Boston, MA.

Vives, X. 1999. *Oligopoly Pricing: Old Ideas and New Tools*. The MIT Press, Cambridge, MS.

Accepted by Wallace J. Hopp; received February 10, 1999. This paper was with the authors 4 months for 2 revisions.