Available online at www.sciencedirect.com
SCIENCE DIRECT®
European Journal of Operational Research 163 (2005) 653–667
EUROPEAN JOURNAL OF OPERATIONAL RESEARCH
www.elsevier.com/locate/dsw

Optimal shared-savings contracts in supply chains: Linear contracts and double moral hazard

Charles J. Corbett [a,*], Gregory A. DeCroix [b], Albert Y. Ha [c]

[a] The Anderson School at UCLA, 110 Westwood Plaza, Box 951481, Los Angeles, CA 90095-1481, USA
[b] Fuqua School of Business, Duke University, Durham, NC 27708-0120, USA
[c] Department of Information & Systems Management, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong

Available online 17 March 2004

Abstract

In many supply chains consumption of indirect materials, sold by a supplier to a customer for use in her production process, can be reduced by efforts exerted by either party. Since traditional supply contracts provide no incentive for the supplier to exert such effort, shared-savings contracts have been proposed as a way to improve incentives in the channel, leading to more efficient effort choices by the two parties. Such shared-savings contracts typically combine a fixed service fee with a variable component based on consumption volume. We formalize this situation using the double moral hazard framework, in which both parties decide how much effort to exert by trading off the cost of their effort against the benefits that they will obtain from reduced consumption. We also extend the double moral hazard framework to analyze a broader class of cost-of-effort functions than considered so far, including the linear cost-of-effort functions commonly found in practice. We show that the supplier can still always induce the optimal second-best equilibrium with a linear shared-savings contract. Under this broader class of functions, however, the behavior of the optimal contract as a function of the problem parameters becomes more complex. We illustrate how small changes in the problem parameters can turn profits from being a well-behaved to a poorly-behaved function of the contract, and provide some theoretical characterization of this phenomenon. The practical significance of this is that simple (linear) contracts are sufficient in many double moral hazard contexts, even for the broader class of functions we consider, but care must be taken in selecting the optimal contract parameters.

Keywords: Supply chain management; Supply contracts; Shared savings; Game theory; Double moral hazard; Cost-of-effort functions; Environmental management; Indirect materials

1. Introduction

In traditional supplier–customer relationships, the supplier's profits are increasing in the quantity of materials consumed by the customer. For some materials the quantity consumed is strictly determined by demand for the customer's product

* Corresponding author. Tel.: +1-310-825-1651; fax: +1-310-206-3337.
E-mail addresses: charles.corbett@anderson.ucla.edu (C.J. Corbett), decroix@mail.duke.edu (G.A. DeCroix), imayha@ust.hk (A.Y. Ha).

0377-2217/$ - see front matter © 2004 Elsevier B.V. All rights reserved.
doi:10.1016/j.ejor.2004.01.021

(e.g., a clothing retailer "consumes" one shirt for every unit sold, a car manufacturer consumes four tires for every car manufactured and sold, etc.). For other materials, however, the link between the customer's sales and consumption is indirect at best. This is true, by definition, for indirect materials such as paint, solvents, office supplies, etc. Both the customer and the supplier can exert effort to reduce the volume the customer consumes of these materials for a given level of customer sales.

Whenever the supplier earns a positive margin on the indirect materials sold to the customer, as is true under traditional contracts, supplier profits are still increasing in the quantity consumed by the customer. From the customer's perspective, though, lower consumption means lower costs. The supplier has no incentive to help the customer reduce consumption, which may result in a sub-optimal outcome for the supply chain. In recent years some companies have addressed this conflict through the use of *shared-savings contracts*. These two-part contracts include a fixed (service) fee and a variable component which is less than the supplier's unit cost of production. Under such contracts both parties have an incentive to reduce consumption. As a result, when compared to a traditional contract, a shared-savings contract generally leads to savings which then need to be shared appropriately; hence the name. Examples of such contracts in practice are discussed in Bierma and Waterstraat (1996, 2000) and Reiskin et al. (2000).

Corbett and DeCroix (2001) analyze the impact of shared-savings contracts on channel profits and material consumption assuming that the variable component is linear in the quantity of indirect material used. While such linear contracts can yield higher profits as well as lower consumption, the resulting equilibrium effort levels are generally not the first-best (i.e., channel-optimal) outcome. This leaves open the possibility that more general (nonlinear) contracts could yield higher profits, possibly even achieving the first-best outcome. This paper addresses the question of determining the *optimal* shared-savings contract (from the supplier's perspective) in a more general setting than that studied in Corbett and DeCroix (2001).

To answer these questions we use the concept of moral hazard. Single (or one-sided) moral hazard arises when a decision-maker does not fully internalize the costs and benefits resulting from his decisions, and those decisions are unobservable or at least unverifiable, so they cannot be contracted on; see, for instance, Van Ackere (1993). Double (or two-sided) moral hazard occurs when the same applies to both parties in a transaction. If decisions are verifiable, the parties can simply contract directly on those decisions, guaranteeing whatever outcome is most desirable. Typically, decisions such as how much effort to exert are not verifiable, so the parties must contract on other factors, such as consumption of indirect materials, that are verifiable. We discuss specific literature on double moral hazard in the next section; a key result there is that the supplier's optimal contract consists of a fixed part and a linear variable part.

Since the shift to shared-savings contracts in practice is often driven by the supplier, we take the supplier's perspective here rather than the channel perspective adopted in Corbett and DeCroix (2001). This paper also generalizes the setting studied in that earlier paper by including a stochastic relationship between effort and consumption, and by allowing for a larger class of cost-of-effort functions. A standard assumption in the double moral hazard literature is that "the first marginal unit of effort is costless". This assumption allows one to rule out boundary equilibria, which facilitates the analysis. However, despite seeming innocent, this assumption rules out linear cost-of-effort functions. In our discussions with practitioners, effort usually consisted of engineering hours, which gives rise to precisely such linear cost-of-effort functions. Analysis of this more realistic setting therefore requires an extension of existing double moral hazard results to include cases with boundary equilibria.

We show that, even in this more general setting, the supplier can still always implement his optimal outcome through a contract with a variable component that is linear in consumption. (This supplier-optimal outcome is referred to as the second-best outcome.) This result adds theoretical justification to our earlier empirical basis for studying linear contracts. We also show that, with the broader class of cost-of-effort functions allowed here, it is sometimes possible for the supplier to implement the first-best outcome. This is never possible with the restricted class of cost-of-effort functions studied in the traditional double moral hazard literature. In addition, in our more general setting, the behavior of the supplier's optimal contract as a function of the problem parameters becomes more complex. Most existing double moral hazard literature is either not concerned with actual computation (as opposed to existence or form) of the optimal contract, or (as in Bhattacharyya and Lafontaine, 1995) find the optimal contract by setting the first derivative of profits with respect to the contract parameter equal to zero. This latter approach is not sufficient in the more general case studied here. We find, numerically, that channel profits are a "W"-shaped function of the contract parameter, and analytically verify part of that shape. In short, this work brings good news and bad news: the good news is that even in our more general setting, there is still always a linear optimal contract, which is therefore easy to implement. The bad news is that the optimal contract parameters are complex functions of the problem parameters and may be difficult to find analytically as first-order conditions are not sufficient (but the interval over which we need to search is bounded).

Finally, the current paper demonstrates the usefulness of the double moral hazard framework in studying supply chains, while also illustrating the limitations of the type of assumption typically made in the economics literature when transplanted to an operations management setting. We hope that this research will prompt additional work along these lines.

In Section 2, we review pertinent literature on supply-chain coordination and on double moral hazard models. Section 3 describes our modeling framework, and Section 4 derives the main results about linearity of optimal contracts and circumstances under which the supplier can achieve first-best. Section 5 contains numerical results for situations which are not well-behaved, and these are theoretically characterized in Section 6. Section 7 concludes the paper.

2. Literature

In this paper we use the double moral hazard framework to analyze settings where shared-savings contracts are applied to transactions involving indirect materials. Bierma and Waterstraat (1996, 2000) and Reiskin et al. (2000) describe various contracts currently used for procurement of chemicals, as formalized in Corbett and DeCroix (2001), discussed earlier.

Holmström (1979) analyzed the single moral hazard problem between a risk-neutral principal and a risk-averse agent. In our setting (with a risk-neutral agent), this would correspond to the case where only the customer can influence consumption. In this case the supplier's optimal (though unconventional) contract would make the customer bear all variable costs of consumption in exchange for a fixed fee, essentially selling the supplier's firm to the customer but extracting all excess channel profits through the selling price. When the supplier can also affect consumption, for instance by sending engineers to the customer's site to suggest process improvements, selling the firm to the customer no longer solves the problem.

Early uses of the double moral hazard framework were in the context of agricultural economics (Reid, 1973) and franchise contracts (Mathewson and Winter, 1985). Cooper and Ross (1985) and Mann and Wissink (1990) study double moral hazard in the context of product warranties. Demski and Sappington (1991) show that the principal can achieve the first-best outcome in the unlikely event that he can observe the agent's effort and can force the agent to buy the firm at a predetermined price if the agent's effort is too low. Gupta and Romano (1998) find that in some instances, the principal can overcome the double moral hazard problem by linking multiple agents' rewards. Bhattacharyya and Lafontaine (1995) and Kim and Wang (1998) show that, in general, first-best cannot be achieved. Under their more restrictive assumption on cost-of-effort functions, Bhattacharyya and Lafontaine (1995) show that linear contracts are optimal and provide empirical evidence that they perform well in franchising contexts; Kim and Wang (1998) show that optimality no longer holds with a risk-averse agent. In this paper we examine to what extent these results can be generalized to a broader class of cost-of-effort functions.

Several recent streams of work in the operations management literature study situations where the double moral hazard framework might be useful. For instance, Cachon and Larivière (in press) show that revenue sharing (such as between videotape distributors and rental outlets) can lead to better results than traditional coordination schemes. Cachon and Zipkin (1999) use game theory to show how cooperation with respect to inventory policies between supplier and retailer could sometimes lead to dramatic improvements in system performance. Caldentey and Wein (2003) show analogous findings when the supplier is modeled as an M/M/1 queue and the retailer determines inventory levels. Dudek (2002) studies negotiation-based collaboration between a buyer and supplier who both use mathematical programming for production planning decisions. Chen (1999) and Lee and Whang (1999) show how designing appropriate incentive mechanisms and information flow can achieve coordination in decentralized supply chains. Porteus (2000) proposes "responsibility tokens" as a way of operationalizing Lee and Whang's incentive scheme. Tsay (1999) examines how appropriate incentives can help achieve coordination when the customer does not commit to a specific order quantity but only to a range. Van Mieghem (1999) examines under which circumstances contracts between a manufacturer and subcontractor should be specified in more or less detail. Porteus and Whang (1991) derive the optimal incentive scheme for a principal facing moral hazard among marketing and manufacturing managers. In Gilbert and Cvsa (2003), a firm has an incentive to underinvest in cost reduction, to increase its own prices; precommitting to a given price may help but reduces the flexibility to respond to demand uncertainty. Gupta and Loulou (1998) study how substitutability changes the effects of inserting retailers between manufacturers and final customers on the manufacturers' profits and investments in process innovation.

Baiman et al. (2000) is the only paper in the operations management literature that we are aware of that explicitly refers to double moral hazard; there, supplier and customer can invest in improving product quality and appraisal quality, respectively, both reducing external failures. They show that first-best can be achieved in scenarios when certain decisions and outcomes are contractible, and then explore performance of the second-best solution in the absence of such contractibility. The type of contracts they consider is different from ours, as their main outcome variable (external failures) is binary rather than continuous.

3. Model and assumptions

The modeling framework we use embeds a generalized version of that introduced in Corbett and DeCroix (2001) in the double moral hazard framework as used by Kim and Wang (1998), but extends the latter's analysis to a broader class of cost-of-effort functions under which effort equilibria need no longer be interior.

A customer consumes Y units of indirect materials purchased from a supplier, where units are scaled so that Y lies between 0 and 1. Both supplier and customer can exert effort e_s and e_c, respectively, to reduce that consumption and the aggregate impact of those efforts is captured by a "composite effort function" $\phi(e_s, e_c)$. Let $\phi(e_s, e_c)$ be continuously differentiable on $(e_s, e_c) \in [0, \infty) \times [0, \infty)$ (using right-hand derivatives at 0), strictly concave and strictly increasing in both parameters. Consumption Y is a random variable that depends stochastically on composite effort through a probability distribution with density $f(y \mid \phi(e_s, e_c))$, where $f(y|z)$ is continuously differentiable in z. We will use superscripts to denote derivatives throughout, so that $f^z(y|z) = \partial f(y|z)/\partial z$, $\phi^s(e_s, e_c) = \partial\phi(e_s, e_c)/\partial e_s$, etc. Assume that $f^z(y|z)$ is bounded on $y \in [0, 1]$ for any given z. Also, we will use the shorthand notation $f^\phi(y \mid \phi)$ to denote $f^z(y|z)$ evaluated at $z = \phi(e_s, e_c)$. Let consumption Y be strictly stochastically decreasing in composite effort $\phi(e_s, e_c)$, so that

$$\frac{\partial E[Y \mid \phi(e_s, e_c)]}{\partial e_s} = \phi^s(e_s, e_c) \int_0^1 y f^\phi(y \mid \phi(e_s, e_c))\, \mathrm{d}y < 0$$

and

$$\frac{\partial E[Y\,|\,\phi(e_s,e_c)]}{\partial e_c} = \phi^c(e_s,e_c)\int_0^1 y f^{\phi}(y\,|\,\phi(e_s,e_c))\,dy < 0.$$

Moreover, assume that expected consumption $E[Y\,|\,\phi(e_s,e_c)]$ is strictly convex in ϕ.

Supplier and customer each incur a cost of exerting effort, given by $c_s(e_s)$ and $c_c(e_c)$. Define $\mathbf{C}$ as the set of cost-of-effort functions $c(e)$ which are continuously differentiable on $[0,\infty)$ and convex increasing in effort, i.e. $dc(e)/de \geqslant 0$ and $d^2c(e)/de^2 \geqslant 0$ for $e\in[0,\infty)$. Define $\mathbf{C}_0 = \mathbf{C}\cap\{c(e)\,|\,dc(0)/de = 0\}$, i.e. the cost-of-effort functions in $\mathbf{C}$ under which the first marginal unit of effort is costless. Here, we only assume $c_s, c_c \in \mathbf{C}$; Corbett and DeCroix (2001), Kim and Wang (1998) and others make the stronger assumption that $c_s, c_c \in \mathbf{C}_0$. This rules out linear cost-of-effort functions such as $c_s(e_s) = k_s e_s$ and $c_c(e_c) = k_c e_c$, which are quite common in practice, as effort often consists primarily of engineering hours. The numerical experiments in Section 5 and theoretical results in Section 6 show how the behavior of the optimal contract becomes more complex once we drop the assumption that $c_s, c_c \in \mathbf{C}_0$.

Unlike Corbett and DeCroix (2001), we take the perspective of the supplier, who can offer a contract requiring the customer to pay an amount $T(y)$, where $T(y)$ is any function of actual observed consumption y. The analysis in Corbett and DeCroix (2001) focused on contracts of the form $T(y) = t + ay$, where t is the fixed fee and a is the unit price. Here, $T(y)$ can be any Lebesgue-integrable function of consumption. Once $T(y)$ has been announced, both parties choose their effort levels simultaneously. The customer earns a fixed revenue r , so that the supplier's and customer's expected profits are

$$\Pi_s(e_s,e_c) = \int_0^1 (T(y) - c\cdot y) f(y\,|\,\phi(e_s,e_c))\,dy - c_s(e_s)$$

and

$$\Pi_c(e_s,e_c) = r - \int_0^1 (T(y) + d\cdot y) f(y\,|\,\phi(e_s,e_c))\,dy - c_c(e_c),$$

where c is the supplier's unit production cost and d is the customer's unit cost for handling and disposal of the indirect material. (In Appendix A we show that these profit expressions are well-defined.) The supplier will choose $T(y)$ so as to maximize his expected profits $\Pi_s(e_s,e_c)$, subject to the requirement that the resulting effort levels are an equilibrium in the effort game for that given contract. The customer will only purchase from the supplier if her expected profits $\Pi_c(e_s,e_c)$ exceed her reservation profit level Π_c^-. All parameters are common knowledge, but effort levels are non-contractible.

4. The suppliers optimal contract

The supplier's optimization problem can be formulated as

$$\begin{aligned}
\mathbf{S}:\quad & \max_{T(y),e_s,e_c} \quad \Pi_s(e_s,e_c) \\
& \text{s.t.} \quad e_s = \arg\max_e \Pi_s(e,e_c), && \text{(ICs)} \\
& \qquad\;\; e_c = \arg\max_e \Pi_c(e_s,e), && \text{(ICc)} \\
& \qquad\;\; \Pi_c(e_s,e_c) \geqslant \Pi_c^-. && \text{(IRc)}
\end{aligned}$$

ICs and ICc are the supplier's and customer's incentive-compatibility constraint, respectively, and define a Nash equilibrium to the effort game. IRc is the customer's individual rationality constraint and ensures her participation. Clearly, the supplier can include a fixed fee in $T(y)$ such that IRc is binding. Hence, adding $\Pi_c(e_s,e_c) = \Pi_c^-$ to the objective function shows that **S** is equivalent to maximizing system profits, subject to ICs and ICc. Any solution to problem **S** is a "second-best" solution. The "first-best" solution (e_s^{**}, e_c^{**}) would arise from solving:

$$\begin{aligned}
\mathbf{FB}:\quad & \max_{T(y),e_s,e_c} \quad \Pi_s(e_s,e_c) \\
& \text{s.t.} \quad \Pi_c(e_s,e_c) \geqslant \Pi_c^-. && \text{(IRc)}
\end{aligned}$$

Here, too, the supplier will set $T(y)$ such that IRc is binding, so **FB** is also equivalent to maximizing system profits, but without the constraints ICs and ICc. **FB** requires that effort levels be contractible, which is generally not true. Note that we include e_c under the maximum operator in **S** and **FB**; though the supplier cannot directly set e_c in $\mathbf{S}$, he can do so indirectly through his choice of $T(y)$.

Our first key result generalizes earlier statements in Bhattacharyya and Lafontaine (1995) and Kim and Wang (1998) to the broader class of cost-of-effort functions allowed here.

Proposition 1. *For any cost-of-effort functions* c_s, $c_c \in \mathbf{C}$, *there exists a linear contract* $T(y) = t^* + a^* y$ *that implements the supplier's second-best outcome, the solution to* $\mathbf{S}$.

The proofs of all results are provided in Appendix A. We include the full details of the proof of Proposition 1 for the generalized case, rather than focus on how it differs from Kim and Wang's (1998) earlier proof. Though Kim and Wang's (1998) proof of optimality of linear contracts for the more restricted class of cost-of-effort functions $\mathbf{C}_0$ was omitted from their paper, they kindly shared it with us. To our knowledge, no rigorous proof has been published yet for c_s, $c_c \in \mathbf{C}_0$, despite several references to the result in Bhattacharyya and Lafontaine (1995), Kim and Wang (1998), and others.

The key steps of the proof are as follows. First, Kim and Wang (1998) show that, for any optimal (possibly nonlinear) contract $T^*(y)$, there exists a linear contract with variable cost a^* which implements the same optimal equilibrium when $c_s, c_c \in \mathbf{C}_0$; in other words, there exists a linear optimal contract. We show that essentially the same a^* (adapted to our setting) is still an optimal linear contract when $a^* \in [-d, c]$, even without this restriction on the cost-of-effort functions. We then show that if $a^* \notin [-d, c]$, we must have a boundary equilibrium; even in that case, a^* still implements the optimal equilibrium. Finally, given any linear contract defined by an $a \notin [-d, c]$, it is always possible to find another linear contract defined by an $a^{**} \in [-d, c]$ in which the supplier's profits are no less than under the previous $a \notin [-d, c]$. Combining all this, we conclude that one can always find an optimal linear contract satisfying $a^{**} \in [-d, c]$, even for any $c_s, c_c \in \mathbf{C}$.

Corbett and DeCroix (2001) show that first-best cannot be achieved with linear contracts under the restriction that $c_s, c_c \in \mathbf{C}_0$. Combined with Proposition 1, this implies that first-best cannot be achieved with any contract for the restricted class of cost-of-effort functions. This is also the case in Bhattacharyya and Lafontaine (1995), Kim and Wang (1998) and elsewhere; the only double moral hazard scenario we are aware of in which the principal can achieve first-best is Demski and Sappington (1991), where the principal can observe the agent's effort (but cannot contract on it) and can force the agent to buy the firm at a predetermined price. However, by considering the broader class of functions $c_s, c_c \in \mathbf{C}$, there are intuitive situations when it *is* possible for the supplier to achieve first-best even under more reasonable assumptions regarding the supplier's knowledge and feasible actions. The following result states this formally and provides sufficient conditions for when first-best can be achieved.

Proposition 2. *In problem* $\mathbf{S}$, *when* $E[Y \mid \phi(e_s, e_c)]$ *is supermodular in* (e_s, e_c), *the supplier can achieve the first-best outcome* (e_s^{**}, e_c^{**}) *as an equilibrium with a linear contract in either of the following situations*:

(*i*) *If* $c_s'(0) \geqslant -(c+d)\frac{\partial E[Y \mid \phi(0,0)]}{\partial e_s}$, *then first-best can be achieved by setting* $a^* = c$ *and* $t^* = r - (c+d)E[Y \mid \phi(0, e_c^{**})] - c_c(e_c^{**}) - \Pi_c^-$.

(*ii*) *If* $c_c'(0) \geqslant -(c+d)\frac{\partial E[Y \mid \phi(0,0)]}{\partial e_c}$, *then first-best can be achieved by setting* $a^* = -d$ *and* $t^* = r - c_c(0) - \Pi_c^-$.

From the proof of Proposition 2, we see that the supplier can achieve first-best when that requires that at most one of the parties exert positive effort. This is reminiscent of the single moral hazard case, in which the principal can achieve first-best by transferring all variable costs to the risk-neutral agent. The crucial distinction with the well-known single moral hazard case is that, in our case, the supplier *chooses*, in a sense, to reduce the double moral hazard problem to a single moral hazard one; this problem is *not* a single moral hazard problem ex ante. The practical significance of this result is that if the supplier is dealing with a severely resource-constrained customer, in which case condition (ii) in Proposition 2 is more likely to hold, it is optimal for the supplier to offer a contract under which the customer has no incentive to exert any effort, i.e., a contract with high fixed fee and a full subsidy for the customer's unit consumption costs. Conversely, if the supplier is the severely resource-constrained party, he should offer a contract combining a relatively low fixed fee with a unit price equal to his own unit cost. If both parties are severely resource-constrained so that conditions (i) and (ii) both hold, then the first-best outcome is $e_s^{**} = e_c^{**} = 0$.

In the preceding analysis we have taken the perspective of the supplier. In some situations, however, it might be more natural for the customer to act as the leader, seeking to maximize her profit by offering to pay an amount $T(y)$ when total consumption is equal to y. Arguments similar to those above can be used to show that customer-oriented versions of Propositions 1 and 2 still hold in this case: a linear contract is sufficient to allow the customer to obtain her second-best profit, and is sufficient to obtain the first-best outcome under the conditions in Proposition 2. Writing Π_s^- for the supplier's reservation profit, the customer-led version of part (i) of Proposition 2 would be: if $c_s'(0) \geqslant -(c+d)\frac{\partial E[Y\,|\,\phi(0,0)]}{\partial e_s}$, then first-best can be achieved by setting $a^* = c$ and $t^* = c_s(0) + \Pi_s^-$. Part (ii) of Proposition 2 becomes: if $c_c'(0) \geqslant -(c+d)\frac{\partial E[Y\,|\,\phi(0,0)]}{\partial e_c}$, then first-best can be achieved by setting $a^* = -d$ and $t^* = (c+d)E[Y\,|\,\phi(e_s^{**},0)] + c_s(e_s^{**}) + \Pi_s^-$.

5. Numerical experiments: Complex behavior of the optimal contract

If the first marginal unit of effort is costless, i.e., $c_s, c_c \in \mathbf{C}_0$ (as in Kim and Wang, 1998 and Bhattacharyya and Lafontaine, 1995), the optimal contract parameter and equilibrium effort levels will be interior. In our setting, this means $-d < a^* < c$ and $e_s^* > 0$, $e_c^* > 0$. This in turn implies that first-order conditions are necessary to characterize such an equilibrium.

Under the more general assumption that $c_s, c_c \in \mathbf{C}$, some linear contracts may lead to an equilibrium with one or both effort levels equal to zero. Total channel profits are then not well-behaved as a function of the contract parameter a, which complicates characterization of the optimal contract and resulting equilibrium. In this section we use numerical examples to illustrate how local maxima emerge when we allow $c_s, c_c \in \mathbf{C}$, and we analytically verify this in Section 6. We experimented with a wide range of scenarios, and selected the examples below as being representative of the behavior we found. In all our experiments, and in our analysis in Section 6, we used consumption functions such that $E[Y\,|\,\phi(e_s, e_c)]$ was supermodular in (e_s, e_c).

Example 1. We start with a variation of an example from Bhattacharyya and Lafontaine (1995, p. 771). Since they focus on the value of joint production while we focus on the cost of consumption, we modify their function slightly. In their example, joint production is $Ke_s^{\beta_s}e_c^{\beta_c}$ with $0 < \beta_s, \beta_c < 1$ (our notation). We assume expected consumption is $E[Y\,|\,\phi(e_s, e_c)] = (1+e_s)^{\beta_s}(1+e_c)^{\beta_c}$ with $\beta_s, \beta_c < 0$; in our examples, $\beta_s = \beta_c = -0.7$. We set $c_s(e_s) = \delta_s(e_s + x_s)^{m_s}/m_s$ and $c_c(e_c) = \delta_c(e_c + x_c)^{m_c}/m_c$ with $\delta_s = \delta_c = 5$ and $m_s = m_c = 2$, so $c_s, c_c \in \mathbf{C}$. When $x_s = x_c = 0$, these functions are of the same form as in Bhattacharyya and Lafontaine (1995), and satisfy $c_s, c_c \in \mathbf{C}_0$. When x_s and x_c are positive, the functions simply shift to the left, so the first marginal unit of effort has positive cost. The other parameters used are $c = d = 10$ and $r = 30$.

Fig. 1 shows the behavior of total channel profits as a function of the contract parameter a for $x_s = x_c = 0$, $x_s = x_c = 1$ and $x_s = x_c = 2$. In the first case, channel profits are concave in a for $-d \leqslant a \leqslant c$, so it is easy to identify the optimal contract as that which shares unit consumption costs equally between the two parties, i.e., $a^* = (c-d)/2 = 0$. (This solution is to be expected given the symmetry in the problem.) The second case, however, exhibits very different behavior. After the change of variable above ($e_s' = e_s + 1$ and $e_c' = e_c + 1$), our example is of the same form as that analyzed by Bhattacharyya and Lafontaine (1995), who use first-order derivatives to find the optimal contract. Naively extending their first-order approach from $c_s, c_c \in \mathbf{C}_0$ to $c_s, c_c \in \mathbf{C}$ would suggest that the optimal contract for this example

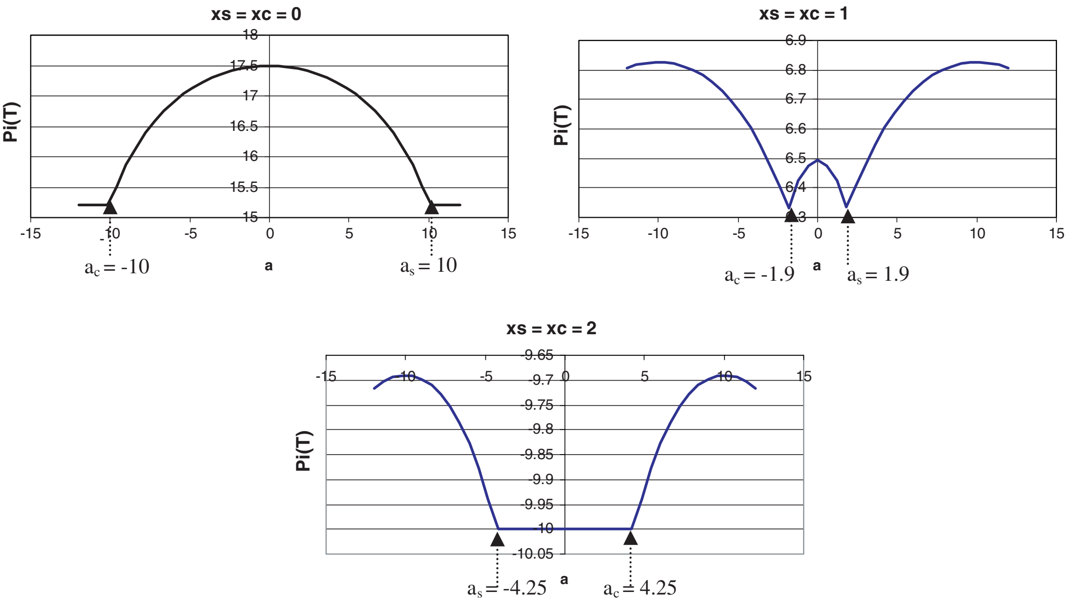

Fig. 1. Channel profits as a function of contract parameter a for Example 1. The variables a_s and a_c are defined in Section 6. These graphs show the impact of dropping the assumptions that $dc_s/de_s(0) = 0$ and $dc_c/de_c(0) = 0$, satisfied only when $x_s = x_c = 0$, on the behavior of channel profits. (Note that the vertical scales for the graphs are different. Due to the forms of the functions used, higher x_s and x_c naturally lead to significantly higher effort costs, which makes direct comparisons of absolute profit levels meaningless. Our focus is on the general shapes of the profit curves.)

is the symmetric one, i.e., $a^* = (c - d)/2 = 0$, and that both parties would exert positive effort in the resulting equilibrium. However, this is only a local maximum for channel profits. The global maximum occurs when $a^* = c = 10$ or $a^* = -d = -10$; in both cases, the party incurring zero unit consumption cost will exert zero effort. The third case, with $x_s = x_c = 2$, does not even have a strict local interior maximum. In Section 6 we show analytically that, whenever $c_s, c_c \in \mathbf{C} \backslash \mathbf{C}_0$ (where $\backslash$ denotes set subtraction), the type of behavior illustrated here will indeed occur.

Example 2. This example illustrates how the same shift from an interior solution to a boundary solution can occur even by changing just a single cost parameter, rather than the cost-of-effort functions as in Example 1. We start with the second case from Example 1 ($x_s = x_c = 1$), and then vary just the handling and disposal cost parameter d, using $d = 20$, 10 and 0. Total channel profits as a function of a are shown in Fig. 2. With $d = 20$, we find a relatively well-behaved scenario with an interior optimal contract ($-d < a^* < c$) and equilibrium ($e_s^* > 0$, $e_c^* > 0$). With $d = 10$, the interior maximum $a^* = (c - d)/2 = 0$ is no longer a global maximum, but instead we have an optimal contract parameter $a^* = c$ or $a^* = -d$ and an equilibrium with one party exerting zero effort. With $d = 0$, we no longer have a strict local interior maximum.

These two examples illustrate two additional insights. First, when the initial marginal cost of effort has some positive cost, the channel profit function tends to take on the "W" shape seen in Figs. 1 and 2. This pattern, partly verified analytically below, recurred consistently in our numerical experiments. If a scenario was not symmetric between the supplier and the customer, then an asymmetric "W" shape was observed. Second, there does not appear to be a simple method for determining a priori whether a given scenario will result in an interior optimal contract or one with

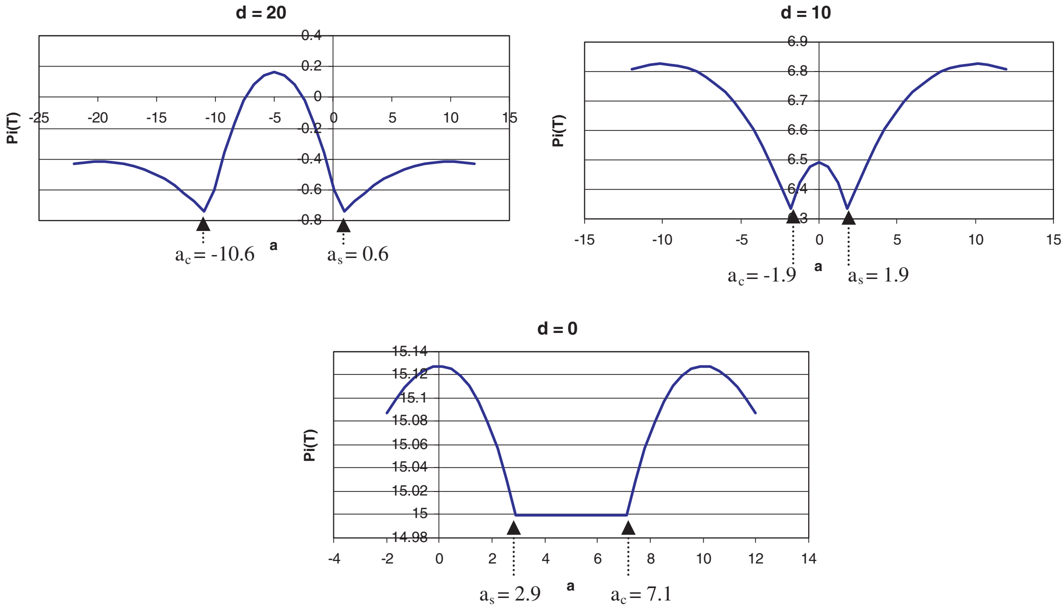

Fig. 2. Channel profits as a function of contract parameter a for Example 2. The variables a_s and a_c are defined in Section 6. This graph shows a set of three scenarios, in which only the parameter d changes, from 20 to 0. (Note again that the horizontal and vertical scales for the graphs are different. Our objective is not to directly compare absolute profit levels among distinct, though similar, scenarios having different handling/disposal costs. Clearly higher costs lead to lower channel profits, and direct performance comparisons are not meaningful. Instead, our focus once again is on the general behavior of the profit curves.)

$a^* = c$ or $a^* = -d$. This further illustrates the risk of relying on first-order conditions to characterize behavior in these kinds of games.

In addition to the examples described above, we also performed similar numerical trials using linear cost-of-effort functions. In all of these trials we observed the same characteristic "W" shape seen in Figs. 1 and 2. We chose to report the results using the nonlinear cost-of-effort functions shown above in order to remain closer to the functions used in existing work, especially Bhattacharyya and Lafontaine (1995).

6. Characterization of profits as function of contract parameter: The "W" shape

In the numerical examples, profits were a "W"-shaped function of a, when $c_s, c_c \in \mathbf{C} \setminus \mathbf{C}_0$. This shape has immediate implications for attempts to find the optimal contract, so we analyze this observation more formally here. Let $E[Y \mid \phi(e_s, e_c)]$ be supermodular in (e_s, e_c), as in Corbett and DeCroix (2001); Kim and Wang (1998) make the slightly stronger assumption that $\phi(e_s, e_c)$ is supermodular in (e_s, e_c), a sufficient but not necessary condition for $E[Y \mid \phi(e_s, e_c)]$ to be supermodular. The profit functions are now submodular in (e_s, e_c), so the comparative statics results from Proposition 2 in Corbett and DeCroix (2001) can be extended to the current situation: if $(e_s^*(a), e_c^*(a))$ is the equilibrium for a given a, then $e_s^*(a)$ decreases and $e_c^*(a)$ increases as a increases. (The formal statement of this result uses the concept of iterated play from Lippman et al. (1987) to allow for multiple equilibria.) These results help establish the following lemma.

Lemma 1. *There exists an interval* $A_c = [-d, a_c)$, *with* $a_c \geqslant -d$, *such that for* $a \in A_c$, $e_c^*(a) = 0$ *and* $e_s^*(a)$ *is decreasing in* a. *Total channel profits decrease in* $a \in A_c$, *and* a_c *is increasing in the quantity* $c_c'(0)$.

The lemma implies that the "left-hand side" of the graph of profits as a function of a will be decreasing for all $a \in A_c$, which is a partial characterization of the "W"-shape observed in the numerical trials. The threshold a_c used to define A_c can be characterized by

$$a_c \equiv \sup\left\{ a \in [-d,c] : \frac{\partial \Pi_c(\hat{e}_s(0;a),0)}{\partial e_c} = -(a+d)\frac{\partial E[Y \mid \phi(\hat{e}_s(0;a),0)]}{\partial e_c} - c_c'(0) \leqslant 0 \right\}. \quad (1)$$

Note that when $c_s, c_c \in \mathbf{C}_0$, we can show that $a_c = -d$, as expected: the left- and rightmost parts of the W-shape vanish (as in the first graph in Fig. 1), leaving only the central part, in which first-order conditions may be sufficient to characterize the global optimum. Symmetric arguments establish analogous results with respect to the supplier, where a_s is characterized by

$$a_s \equiv \inf\left\{ a \in [-d,c] : \frac{\partial \Pi_s(0,\hat{e}_c(0;a))}{\partial e_s} = (a-c)\frac{\partial E[Y \mid \phi(0,\hat{e}_c(0;a))]}{\partial e_s} - c_s'(0) \leqslant 0 \right\}.$$

Lemma 2. *There exists an interval* $A_s = (a_s, c]$, *with* $a_s \leqslant c$, *such that for* $a \in A_s$, $e_s^*(a) = 0$ *and* $e_c^*(a)$ *is increasing in a. Total channel profits increase in* $a \in A_s$, *and* a_s *is decreasing in the quantity* $c_s'(0)$.

Some sensitivity analyses show how the intervals A_s and A_c change with c and d.

Lemma 3. a_c *and* a_s *both increase with c and decrease with d.*

Fig. 2 illustrates Lemma 3; as d decreases, both a_c and a_s increase. We have not been able to characterize the central part of the W-shape for the broader class of cost-of-effort functions $\mathbf{C}$, and we are not aware of any characterization for the more limited class $\mathbf{C}_0$. However, if $a_s \leqslant a_c$, neither party exerts any effort for $a \in (a_s, a_c)$. Channel profits are flat on this region, and the only strict local maxima are $a = -d$ and $a = c$, as in Figs. 1 and 2. If $a_s > a_c$, both parties exert positive effort on (a_c, a_s), but without further assumptions on $c_s(e_s)$, $c_c(e_c)$ and $\phi(e_s, e_s)$ we are not able to characterize channel profits for $a \in (a_c, a_s)$. Define

$$A = \left\{ a \in [-d,c] \,\middle|\, \frac{\mathrm{d}\Pi_s(\hat{e}_s(a),\hat{e}_c(a))}{\mathrm{d}a} + \frac{d\Pi_c(\hat{e}_s(a),\hat{e}_c(a))}{da} = 0 \right\},$$

i.e. A is the set of contract parameters a that satisfy the first-order conditions for optimality. Our (partial) results on the supplier's optimal contract can then be summarized as follows.

Proposition 3. *The variable part of the supplier's optimal contract* a^* *satisfies* $a^* \in A \cup \{-d, c\}$, *i.e.,* a^* *either satisfies the first-order-conditions or* $a^* \in \{-d, c\}$.

In other words, despite the non-concavity of the supplier's profit function once one allows for cost-of-effort functions in $\mathbf{C} \backslash \mathbf{C}_0$, the set of possible optima that would need to be evaluated is quite limited.

7. Conclusions

In this paper we have examined the situation where both supplier and customer can exert effort to reduce consumption of an indirect material. We have extended the double moral hazard literature to allow for a broader class of cost-of-effort functions, including the linear functions found in practice, and have shown that the supplier's optimal contract still consists of a fixed part and a variable part which is linear in consumption. We find that, in some cases, the supplier can achieve the first-best outcome, in contrast to the common finding in the double moral hazard literature with more restricted cost-of-effort functions, though in general he still cannot. We have shown that the optimal contract always lies within a finite range, which makes a simple numerical line search easy to implement. We show that under our broader class of cost-of-effort functions, the first-order approach used elsewhere can lead to incorrect conclusions in determining the supplier's optimal contract. We used numerical experiments to illustrate the sometimes complex relationship between supply chain profits and the contract parameter. We have shown that, for the broader class of cost-of-effort functions, channel profits are a "W"-shaped function of the contract parameter. We summarize this in a partial characterization of the optimal contract, though more work is needed to provide a complete description.

To the best of our knowledge, this is the first time the general double moral hazard framework has been used in the context of supply chain contracting; we believe that many other questions in the supply chain literature would benefit from the double moral hazard perspective. In particular, it can help identify when one may or may not be able to achieve first-best, and when a linear (and hence easily implementable) contract may or may not be optimal.

Possible future work includes extending Kim and Wang's (1998) analysis of a risk-averse agent to our case with linear and other cost-of-effort functions, to verify whether the linear contract's lack of robustness to risk aversion still occurs. Another extension would let a single supplier contract with multiple customers (or vice versa), allowing for various types of interaction between them. Finally, a more extensive empirical analysis of the shape and effects of supply contracts in practice is much needed.

Acknowledgements

The authors are grateful to Sugato Bhattacharyya and Susheng Wang for sharing their work with us, in particular the proofs that were omitted from their published papers. The research of Albert Ha was partly supported by HKUST grant DAG 01/02.BM35.

Appendix A

Proof of Proposition 1. First, to see that Π_s and Π_c are well-defined for any optimal contract $T(y)$, note that since $T(y)$ is Lebesgue-integrable, it must be finite almost everywhere. Given our assumptions on existence and continuous differentiability of $f(y\,|\,\phi)$, one can replace $T(y)$ with a $\widetilde{T}(y)$ which is finite everywhere without changing expected profits. Therefore, without loss of generality, we can restrict attention to $T(y)$ finite everywhere. Since $T(y)$ and $f(y\,|\,\phi)$ are both Lebesgue-integrable, they are also Lebesgue-measurable, so, by Theorem 10.36 in Apostol (1982), $T(y)f(y\,|\,\phi)$ is also Lebesgue-measurable. If $T(y)f(y\,|\,\phi)$ is not Lebesgue-integrable, we must have either $\int_0^1 T(y) f(y\,|\,\phi)\,\mathrm{d}y \in \{-\infty, +\infty\}$, in which case also either $\Pi_s = -\infty$ and $\Pi_c = \infty$ or vice versa, or $\int_0^1 T(y) f(y\,|\,\phi)\,\mathrm{d}y$ is not well-defined. Clearly, in both cases, the contract $T(y)$ is not optimal or not feasible, so we can restrict attention to $T(y)$ such that $T(y)f(y\,|\,\phi)$ is Lebesgue-integrable, which means that Π_s and Π_c are well-defined for any optimal $T(y)$.

Since $T(y)$ can include a fixed payment term, the supplier can always guarantee that IRc in **S** holds with equality. Adding IRc to the objective function of **S** (and then dropping the constant term Π_c^-) yields the modified supplier's problem:

$$\textbf{MS}:\quad \max_{T(y),e_s,e_c} \quad \Pi_T(e_s,e_c) := \Pi_s(e_s,e_c) + \Pi_c(e_s,e_c)$$
$$\text{s.t.}\quad \text{ICs, ICc.}$$

Now let $(T^*(\cdot), e_s^*, e_c^*))$ be a solution to **MS** – the equilibrium (e_s^*, e_c^*) may be an interior solution, or either or both of the effort levels may be zero. We assume throughout that there is a unique effort equilibrium that implements that second-best outcome; if there are multiple equilibria that implement second-best, we assume the supplier can enforce any particular one of those equilibria. We want to show that the supplier can still achieve second-best even if he restricts attention to linear contracts $T(y) = t + ay$. In other words, ignoring the fixed fee t, the solution (a^*, e_s^*, e_c^*) to the linear modified supplier problem (**LMS**) below also solves **MS**.

$$\textbf{LMS}:\quad \max_{a,e_s,e_c} \quad \Pi_T(e_s,e_c)$$
$$\text{s.t.}\quad \text{ICs, ICc.}$$

There are several cases to consider. Below, we first consider the case that $e_s^* = e_c^* = 0$, after which we consider the case that $e_s^* > 0$ and e_c^* is either zero or positive. The final remaining case, where $e_s^* = 0$ and $e_c^* > 0$, can then be handled by analogous arguments.

First, suppose that the second-best effort levels are $e_s^* = e_c^* = 0$. In that case, if the supplier offers a simple linear contract with $a = -d$ and t set appropriately, the customer will always choose $e_c^* = 0$, which is the second-best effort level, regardless of e_s. (This follows from a slight modification of Proposition 1 in Corbett and DeCroix, 2001.) Now with $a = -d$, the supplier bears the full channel costs related to consumption, so he will choose the channel-optimal response to $e_c^* = 0$, i.e., he will choose $e_s^* = 0$. To see this, suppose instead that the supplier's best response is $\hat{e}_s > 0$, and that this choice is strictly better than $e_s^* = 0$. This would imply

$$\begin{aligned}\Pi_s(\hat{e}_s, 0) &= t - (c+d)E[Y \mid \phi(\hat{e}_s, 0)] - c_s(\hat{e}_s) \\ &> t - (c+d)E[Y \mid \phi(0,0)] - c_s(0) \\ &= \Pi_s(0,0),\end{aligned}$$

and so

$$\begin{aligned}&r - (c+d)E[Y \mid \phi(\hat{e}_s, 0)] - c_s(\hat{e}_s) - c_c(0) \\ &\quad > r - (c+d)E[Y \mid \phi(0,0)] - c_s(0) - c_c(0),\end{aligned}$$

i.e., $\Pi_T(\hat{e}_s, 0) > \Pi_T(0,0)$. Therefore $(\hat{e}_s, 0)$ would be an equilibrium yielding strictly higher channel profits than the second-best solution $(e_s^*, e_c^*) = (0,0)$, a contradiction. The supplier's optimal response to $e_c^* = 0$ under $a = -d$ must be $e_s^* = 0$, so second-best is achieved under that contract.

Now consider the case where $e_s^* > 0$ and $e_c^* \geqslant 0$. Since $T(y) - cy$ is Lebesgue-integrable and $f^\phi(y \mid \phi(e_s, e_c))$ is bounded on $y \in [0,1]$, $[T(y) - cy] f^\phi(y \mid \phi(e_s, e_c))$ is Lebesgue-integrable and therefore, by Theorem 10.39 of Apostol (1982), we can differentiate Π_s under the integral sign to get

$$\begin{aligned}\partial \Pi_s / \partial e_s(e_s, e_c) = \phi^s(e_s, e_c) \int_0^1 [T(y) - cy] f^\phi \\ \times (y \mid \phi(e_s, e_c))\, \mathrm{d}y - c_s'(e_s).\end{aligned}$$

As $c_s(e_s)$, $\phi(e_s, e_c)$ and $f(y \mid \phi)$ are continuously differentiable by assumption, and since by Theorem 10.38 of Apostol (1982) the first term of $\partial \Pi_s / \partial e_s(e_s, e_c)$ is continuous, Π_s is continuously differentiable in e_s. Likewise, Π_c is continuously differentiable in e_c. Then the following conditions are *necessary* for the second-best outcome (e_s^*, e_c^*) to be an equilibrium:

$$\frac{\partial \Pi_s}{\partial e_s}(e_s^*, e_c^*) = 0,$$

$$\frac{\partial \Pi_c}{\partial e_c}(e_s^*, e_c^*) = \begin{cases} 0 & \text{if } e_c^* > 0, \\ \leqslant 0 & \text{if } e_c^* = 0. \end{cases}$$

These modified first-order conditions can be written as

$$\phi^s(e_s^*, e_c^*) \int_0^1 [T^*(y) - cy] f^\phi(y \mid \phi(e_s^*, e_c^*))\, \mathrm{d}y = c_s'(e_s^*), \tag{A.1}$$

$$\begin{aligned}&\phi^c(e_s^*, e_c^*) \int_0^1 [T^*(y) + \mathrm{d}y] f^\phi(y \mid \phi(e_s^*, e_c^*))\, \mathrm{d}y \\ &\quad = \begin{cases} -c_c'(e_c^*) & \text{if } e_c^* > 0, \\ \geqslant -c_c'(e_c^*) & \text{if } e_c^* = 0. \end{cases}\end{aligned} \tag{A.2}$$

The linear contract we are seeking is analogous to that given by Kim and Wang (1998), i.e.,

$$a^* = \frac{\int_0^1 T^*(y) f^\phi(y \mid \phi(e_s^*, e_c^*))\, \mathrm{d}y}{\int_0^1 y f^\phi(y \mid \phi(e_s^*, e_c^*))\, \mathrm{d}y}$$

which is finite and well-defined because:

- Since $f^\phi(y \mid \phi)$ is bounded on $y \in [0,1]$ for any ϕ, $\int_0^1 y f^\phi(y \mid \phi(e_s^*, e_c^*))\, \mathrm{d}y$ must be finite. Since
  $$\begin{aligned}&\partial E[Y \mid \phi(e_s^*, e_c^*)] / \partial e_s \\ &\quad = \phi^s(e_s^*, e_c^*) \int_0^1 y f^\phi(y \mid \phi(e_s^*, e_c^*))\, \mathrm{d}y < 0\end{aligned}$$
  and $\phi^s(e_s^*, e_c^*) > 0$ by our assumptions, this implies that the denominator $\int_0^1 y f^\phi(y \mid \phi(e_s^*, e_c^*))\, \mathrm{d}y < 0$.
- Since $f^\phi(y \mid \phi)$ is bounded and $T^*(y)$ is Lebesgue-integrable, $\int_0^1 T^*(y) f^\phi(y \mid \phi(e_s^*, e_c^*))\, \mathrm{d}y$ is finite.

With the contract $T(y) = t^* + a^* y$ with a^* as above, Eqs. (A.1) and (A.2) imply

$$\phi^s(e_s^*, e_c^*)(a^* - c) \int_0^1 y f^\phi(y \mid \phi(e_s^*, e_c^*))\, \mathrm{d}y = c_s'(e_s^*), \tag{A.1L}$$

$$\phi^{c}(e_s^*, e_c^*)(a^* + d)\int_0^1 y f^{\phi}(y\,|\,\phi(e_s^*, e_c^*))\,dy = \begin{cases} -c_c'(e_c^*) & \text{if } e_c^* > 0, \\ \geqslant -c_c'(e_c^*) & \text{if } e_c^* = 0. \end{cases} \tag{A.2L}$$

Eqs. (A.1L) and (A.2L) could be interpreted as the modified first-order conditions for **LMS**. We distinguish two cases: either $a^* \in [-d, c]$ or $a^* \notin [-d, c]$. We first show that, in both cases, supplier and buyer will optimally choose the second-best effort levels (e_s^*, e_c^*), after which we confirm that a^* is therefore an optimal contract for the supplier.

Faced with $a^* \in [-d, c]$, supplier and customer will optimally choose (e_s^*, e_c^*), because both profit functions are concave, so conditions (A.1L) and (A.2L) are sufficient for optimality.

Suppose $a^* \notin [-d, c]$; the equilibrium (e_s^*, e_c^*) is still achieved, but the argument is slightly less direct. Since $e_s^* > 0$, a slightly modified version of Proposition 1 in Corbett and DeCroix (2001) establishes that $a^* < c$, which then implies $a^* < -d$. Assume that $e_c^* > 0$, so that (A.2) must be satisfied with equality. $\phi^c(e_s^*, e_c^*)$ is finite, so from the definition of a^* and Eq. (A.2) we have

$$a^* = \frac{\int_0^1 T^*(y) f^{\phi}(y\,|\,\phi(e_s^*, e_c^*))\,dy}{\int_0^1 y f^{\phi}(y\,|\,\phi(e_s^*, e_c^*))\,dy} = -d + \frac{-c_c'(e_c^*)}{\phi^c(e_s^*, e_c^*)\int_0^1 y f^{\phi}(y\,|\,\phi(e_s^*, e_c^*))\,dy}.$$

The assumptions imply that the rightmost term is non-negative, or $a^* \geqslant -d$, a contradiction, so we must have $e_c^* = 0$. Under any linear contract with $a^* < -d$, the customer will optimally choose $e_c = 0$, the second-best effort level in this case. The supplier's profit function is still concave in e_s, so from (A.1L) choosing $e_s = e_s^*$ is optimal for the supplier too. Summarizing: when $a^* < -d$, it is still optimal for both parties to choose the second-best effort levels (e_s^*, e_c^*). The case where $e_c^* > 0$ and $e_s^* = 0$ is analogous.

We have shown that (a^*, e_s^*, e_c^*) always solves **LMS**. Since (a^*, e_s^*, e_c^*) implements the second-best outcome, the solution to **MS**, we have shown that a linear contract is still optimal even for the broader class **C** of cost-of-effort functions allowed here.

We still need to show that, given any linear contract with $a \notin [-d, c]$, it is always possible to find another linear contract with $a^{**} \in [-d, c]$ in which the supplier's profits are at least as high. To do so, we modify arguments in Corbett and DeCroix (2001) (Proposition 1). Consider $a \geqslant c$; for any such contract $e_s^*(a) = 0$. Since

$$\partial^2 \Pi_c / \partial a \partial e_c(0, e_c) = -\phi^c(0, e_c)\int_0^1 y f^{\phi}(y\,|\,\phi(0, e_c))\,dy \geqslant 0,$$

we have $de_c^*(a)/da \geqslant 0$. Also, the modified first-order condition for e_c must be satisfied, i.e.,

$$-\phi^c(0, e_c^*(a))(a + d)\int_0^1 y f^{\phi}(y\,|\,\phi(0, e_c^*(a)))\,dy - c_c'(e_c^*(a)) \leqslant 0.$$

Since $a \geqslant c$ and $\phi^c \int_0^1 y f^{\phi}\,dy < 0$, we have

$$\frac{\partial \Pi_T}{\partial a} = \left[-\phi^c(0, e_c^*(a))(c + d) \times \int_0^1 y f^{\phi}(y\,|\,\phi(0, e_c^*(a)))\,dy - c_c'(e_c^*(a))\right] \times \left(\frac{de_c^*(a)}{da}\right)$$
$$\leqslant \left[-\phi^c(0, e_c^*(a))(a + d) \times \int_0^1 y f^{\phi}(y\,|\,\phi(0, e_c^*(a)))\,dy - c_c'(e_c^*(a))\right] \times \left(\frac{de_c^*(a)}{da}\right) \leqslant 0.$$

So $\frac{\partial \Pi_T}{\partial a} \leqslant 0$, i.e., as long as $a \geqslant c$, channel profits (weakly) decrease as a increases. Since **LMS** seeks to maximize Π_T, it follows that it is always optimal to set $a^{**} \leqslant c$. An analogous argument shows that $a^{**} \geqslant -d$.

To conclude, we have demonstrated that the supplier can always find a linear optimal contract such that $a^{**} \in [-d, c]$, regardless of whether the second-best solution is achieved at an interior or a boundary effort equilibrium. □

Proof of Proposition 2. If $c_s'(0) \geqslant -(c + d)\partial E[Y\,|\,\phi(0, 0)]/\partial e_s$, then the supermodularity of $E[Y\,|\,\phi(e_s, e_c)]$ implies that $c_s'(0) \geqslant -(c + d)\partial E[Y\,|\,\phi(0,$

$e_c)]/\partial e_s$ for any e_c, i.e., $\partial\Pi_T/\partial e_s(0,e_c) \leqslant 0$ for any e_c, so the first-best outcome has $e_s = 0$. If the supplier sets $a = c$, then both the supplier's and customer's profit functions are concave. Since $\partial\Pi_s(e_s,e_c)/\partial e_s = -c_s'(e_s) \leqslant 0$, the supplier's equilibrium effort will also be $e_s = 0$ regardless of the customer's effort. Also under this contract

$$\frac{\partial\Pi_c(0,e_c)}{\partial e_c} = -(c+d)\frac{\partial E[Y\,|\,\phi(0,e_c)]}{\partial e_c} - c_c'(e_c) = \frac{\partial\Pi_T(0,e_c)}{\partial e_c}$$

so the customer's equilibrium effort level will match first-best customer effort. By setting t^* as suggested the supplier captures all profit beyond the customer's reservation profit. A symmetric argument establishes part (ii) of the result. □

Proof of Lemma 1. Suppose $\hat{e}_c(\hat{e}_s(0;a);a) = 0$ for some $a \in [-d,c]$, where $\hat{e}_i(e;a)$ denotes player i's best response function given contract a. Then $e_c^*(a) = 0$, so $(e_s^*(a),0)$ is an equilibrium, where $e_s^*(a) = \hat{e}_s(0;a)$. For $a \in [-d,c]$, each player's profit function is concave in effort, so $\hat{e}_c(\hat{e}_s(0;a);a) = 0$ if and only if

$$\frac{\partial\Pi_c(\hat{e}_s(0;a),0)}{\partial e_c} = -(a+d)\frac{\partial E[Y\,|\,\phi(\hat{e}_s(0;a),0)]}{\partial e_c} - c_c'(0) \leqslant 0.$$

When $a = -d$ we have $\partial\Pi_c/\partial e_c = -c_c'(e_c) \leqslant 0$, so $e_c^*(-d) = 0$ and $(e_s^*(-d),0)$ is an equilibrium. Supplier profits are submodular in (e_s,a) so $\hat{e}_s(0;a)$ is decreasing in a, which together with the supermodularity of $E[Y\,|\,\phi(e_s,e_c)]$ in (e_s,e_c) implies that $-(a+d)\partial E[Y\,|\,\phi(\hat{e}_s(0;a),0)]/\partial e_c$ is increasing in a. As a result $e_c^*(a) = 0$ for $a \in A_c = [-d,a_c)$, where a_c is given in (1). This also shows that $e_s^*(a) = \hat{e}_s(0;a)$ is decreasing in a for $a \in A_c$.

The preceding fact implies that if $\hat{e}_s(0;-d) = 0$, then $\hat{e}_s(0;a) = 0$ for all $a \in A_c$. For values of $a \in A_c$ where $\hat{e}_s(0;a) = 0$, channel profits are constant at $\Pi_T(0,0) = r - (c+d)E[Y\,|\,\phi(0,0)] - c_s(0) - c_c(0)$. Suppose instead that there exist contracts $a \in A_c$ such that $\hat{e}_s(0;a) > 0$. For such contracts, the first-order condition

$$\frac{\partial\Pi_s}{\partial e_s} = -(c-a)\frac{\partial E[Y\,|\,\phi(e_s,0)]}{\partial e_s} - c_s'(e_s) = 0 \quad \text{(A.3)}$$

is sufficient for the supplier's best response since Π_s is concave in e_s. Now consider the supplier effort e_s^T that maximizes total channel profits Π_T, characterized by

$$\frac{\partial\Pi_T}{\partial e_s} = -(c+d)\frac{\partial E[Y\,|\,\phi(e_s^T,0)]}{\partial e_s} - c_s'(e_s^T) = 0 \quad \text{(A.4)}$$

since Π_T is concave in e_s, which also implies that Π_T is decreasing as e_s increases or decreases from e_s^T. By setting $a = -d$, (A.3) becomes identical to (A.4), so $\hat{e}_s(0;-d) = e_s^T$; the supplier's best response under that contract is the effort level that maximizes channel profit. Since the only impact on channel profits of increasing a on A_c is to decrease supplier effort, total channel profits decrease in a on A_c. The last part of the proposition follows immediately from the characterization of a_c in (1). □

Proof of Lemma 3. The only effect of c on a_c is through $\partial E[Y\,|\,\phi(\hat{e}_s(0;a),0)]/\partial e_c$. Since Π_s is supermodular in (e_s,c) we know that $\hat{e}_s(0;a)$ is increasing in c, which combined with the supermodularity of $E[Y\,|\,\phi(e_s,e_c)]$ in (e_s,e_c) implies that $\partial E[Y\,|\,\phi(\hat{e}_s(0;a),0)]/\partial e_c$ is increasing in c. Since we only consider $a \geqslant -d$, the left-hand side of the inequality in (1) is decreasing in c, so a_c is increasing in c. As d increases, the only impact on (1) is through the coefficient $(a+d)$. Since by assumption $-\partial E[Y\,|\,\phi(\hat{e}_s(0;a),0)]/\partial e_c \geqslant 0$, the left-hand side of the inequality in (1) is increasing in d, so a_c is decreasing in d. Similar arguments apply to a_s. □

References

Apostol, T.M., 1982. Mathematical Analysis. Addison-Wesley, Reading, MA.

Baiman, S., Fischer, P.E., Rajan, M.V., 2000. Information, contracting, and quality costs. Management Science 46 (6), 776–789.

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

Bierma, T.J., Waterstraat, F.L., 1996. P2 Assistance from Your Supplier, Pollution Prevention Review (Autumn), pp. 13–24.

Bierma, T.J., Waterstraat Jr., F.L., 2000. Chemical Management: Reducing Waste and Cost Through Innovative Supply Strategies. John Wiley, New York.

Cachon, G.P., Lariviѐre, M.A., in press. Supply chain coordination with revenue sharing contracts: Strengths and limitations. Management Science.

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

Caldentey, R., Wein, L.M., 2003. Analysis of a decentralized production–inventory system. Manufacturing & Services Operations Management 5 (1), 1–17.

Chen, F., 1999. Decentralized supply chains subject to information delays. Management Science 5 (8), 1076–1090.

Cooper, R., Ross, T.W., 1985. Product warranties and double moral hazard. RAND Journal of Economics 16, 103–113.

Corbett, C.J., DeCroix, G.A., 2001. Shared savings contracts for indirect materials in supply chains: Channel profits and environmental impacts. Management Science 47 (7), 881–893.

Demski, J.S., Sappington, D.E.M., 1991. Resolving double moral hazard problems with buyout agreements. RAND Journal of Economics 22 (2), 232–240.

Dudek, G., 2002. Negotiation-based collaborative planning in supply chains. Working Paper, Darmstadt University of Technology.

Gilbert, S.M., Cvsa, V., 2003. Strategic commitment to price to stimulate downstream innovation in a supply chain. European Journal of Operational Research 150 (3), 617–639.

Gupta, S., Romano, R.E., 1998. Monitoring the principal with multiple agents. RAND Journal of Economics 29 (2), 427–442.

Gupta, S., Loulou, R., 1998. Process innovation, product differentiation, and channel structure: Strategic incentives in a duopoloy. Marketing Science 17 (4), 301–316.

Holmström, B., 1979. Moral hazard and observability. The Bell Journal of Economics 10, 74–91.

Kim, S.K., Wang, S., 1998. Linear contracts and the double moral hazard. Journal of Economic Theory 82, 342–378.

Lee, H., Whang, S., 1999. Decentralized multi-echelon supply chains: Incentives and information. Management Science 45 (5), 633–640.

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

Mann, D.P., Wissink, J.P., 1990. Hidden actions and hidden characteristics in warranty markets. International Journal of Industrial Organization 8, 53–71.

Mathewson, G.F., Winter, R., 1985. The economics of franchise contracts. Journal of Law and Economics 28, 503–526.

Porteus, E.L., 2000. Responsibility tokens in supply chain management. Manufacturing and Service Operations Management 2 (2), 203–219.

Porteus, E., Whang, S., 1991. On manufacturing/marketing incentives. Management Science 37 (9), 1166–1181.

Reid Jr., J.D., 1973. Sharecropping as an understandable market response: The post-Bellum South. Journal of Economic History (March), 106–130.

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

Tsay, A.A., 1999. The quantity flexibility contract and supplier–customer incentives. Management Science 45 (10), 1339–1358.

Van Ackere, A., 1993. The principal/agent paradigm: Its relevance to various functional fields. European Journal of Operational Research 70, 83–103.

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