R&D Cost Sharing along the Supply Chain

Mark R. Frascatore[*]

Faculty of Economics and Financial Studies, Clarkson University, U.S.A.

Abstract

A model of R&D cost sharing between a manufacturer and its component supplier is examined. The manufacturer can pay for a fraction of the supplier's cost-reducing R&D in return for a lower component price, and both firms can improve profits.

Key words: R&D cooperation; supply chain collaboration; research joint ventures

JEL classifications: L14; L24

1. Introduction

In recent years, businesses have increased their use of both domestic and foreign outsourcing of components. With this trend comes a greater likelihood of logistical difficulties between firms in a supply chain due to language barriers, higher monitoring costs, difficulties in quality control, etc.[1] At the same time and in an effort to improve product quality and production methods, manufacturers and their suppliers are placing increased emphasis on sharing the responsibility for the design and production of components. Furthermore, it is well known that in most countries there are far fewer legal restrictions on joint ventures between vertically-related firms than between horizontally-related firms.[2] This seems to provide some compelling evidence for more and better cooperative relationships between such vertically related businesses. To this end, this paper examines specifically the case of R&D cooperation.

Studies of the potential benefits of R&D cooperation have largely been limited to cooperative arrangements between horizontally-related firms. The seminal work by d'Aspremont and Jacquemin (1988) and Kamien et al. (1992) spawned a large amount of research on this topic, but only a handful of papers have analyzed cooperative research arrangements between vertically-related firms in a supply chain in the event that vertical integration is not efficient.[3]

Among the first to model this arrangement was Banerjee and Lin (2001). They

examine an R&D cost-sharing agreement between one or more manufacturers and one or more of their suppliers, in which a supplier can incur a fixed R&D cost that lowers its marginal cost by a fixed amount. The model examines the optimal research joint venture (RJV) size, as measured by the number of firms participating and sharing R&D costs. Ishii (2004) analyzes cooperative R&D between vertically-related firms with duopoly suppliers and duopoly final-good manufacturers. Each supplier and manufacturer can conduct cost-reducing R&D to lower its own marginal production costs, and there exist both vertical and horizontal spillovers of R&D knowledge. Firms can form either vertical or horizontal R&D cartels in which they jointly determine their R&D expenditures to maximize combined profits. Findings are that the vertical RJV cartels yield the largest social welfare if the spillover rate between the suppliers is not too high.[4]

The present model adds to the current body of work first by relaxing the assumption that the supplier must choose a fixed amount of cost-reducing R&D. By allowing the supplier to choose its level of R&D, it is possible to determine the amount that is optimal for the supplier and the amount that is optimal for the supply chain. Also, the present model considers not a joint determination of each firm's own cost-reducing R&D, but rather the fraction of the supplier's R&D that is paid by the manufacturer. This allows a comparison between the non-cooperative and cooperative levels of R&D cost sharing and provides a practical way to achieve the cooperative level of cost-reducing R&D. Furthermore, a parameter is included to describe the fraction of cost-reducing R&D that is passed along to the manufacturer through a lower component price. This parameter helps provide some insight regarding the conditions, such as the R&D spillover rates and the competitive environments, under which a vertical RJV is most effective.

To make these comparisons, the model considers two cases. The first is the non-cooperative case, in which the manufacturer independently chooses a fraction of the supplier's R&D costs to pay, and it benefits through a lower component price. The second is the cooperative case, in which the firms form an agreement or RJV specifying the fraction of the costs to be paid by the manufacturer and a fixed fee to be paid from the supplier to the manufacturer to ensure participation. Results are that in the non-cooperative case the manufacturer chooses to share too little of the supplier's costs, resulting in a sub-optimal level of cost-reducing R&D. In fact, in many instances, the manufacturer chooses to pay for none of the R&D, even if it would lower its component price. The optimal level of sharing under the cooperative case is determined and found to be equal to the fraction of the cost reduction that is passed along to the manufacturer through the component price. Furthermore, cooperation has the greatest effect when that fraction is not too low or too high. That is, cooperation may be most important for intermediate R&D spillover rates and similar levels of competition in the supplier's and manufacturer's markets.

The paper is arranged as follows: Section 2 starts by providing the setup and timing of the model and outlining the two cases. Then the solutions to the cases are characterized and compared. Section 3 illustrates the results with an example, and Section 4 provides concluding remarks.

2. The Model

The firms in the model are vertically-related businesses in a supply chain: a manufacturer and its supplier. Due to high agency efficiency and/or low technical complementarities, vertical integration is assumed to be undesirable. The supplier produces a custom-made, critical component for the manufacturer, and can conduct R&D to lower its production costs for this particular component. A reduction in the supplier's production costs can, in turn, reduce the price it charges the manufacturer for the components. For example, the amount of the component price reduction would be positively related to the level of competition in the supplier's market and/or the R&D spillover rate, and negatively related to the level of competition in the manufacturer's market.[5] No specific assumptions are made about the combination of these factors; rather, a parameter is included to allow for any fraction of the cost reduction to be passed on through the component price. As a simplifying assumption, the manufacturer's output price is assumed to be fixed.[6]

Before the supplier conducts its R&D, the two firms have an opportunity to form an RJV stipulating a fraction of the R&D to be paid by the manufacturer. The RJV can also stipulate any fixed payment from the supplier to the manufacturer to ensure participation.

The Timing

The model chronology is outlined as follows.
1. For simplicity, the manufacturer's price is assumed to be fixed at p. The quantity demanded for the manufacturer's product, Q, is then realized and becomes common knowledge between the supplier and manufacturer.
2. The manufacturer determines and announces it will pay the fraction α of the supplier's R&D costs. The manufacturer may choose α independently or cooperatively with the supplier through an RJV. The firms may also use an RJV to designate a fixed payment R from the supplier to the manufacturer.
3. The supplier chooses a level of R&D denoted y. The value of y equals the reduction in the supplier's marginal cost of producing the components for the manufacturer.[7] The cost of this level of R&D is denoted $f(y)$, where $f'(y)$ and $f''(y)$ exist.
4. The supplier produces Q components at a constant marginal cost of $c_s - y$ and sells to the manufacturer at a per-unit price of $w - \beta_y$, where $\beta \in [0, 1]$.[8]
5. The manufacturer pays the supplier a sum of $Q(w - \beta y) + \alpha f(y) - R$. This includes payment for the components, a fraction of the supplier's R&D expenditures, and the fixed payment (which is cooperatively determined under the RJV, but is zero otherwise).
6. The manufacturer transforms each component into a unit of the final product at a cost of c_M and sells Q units at price p.

The Profit Measures

With the cost and price information in hand we can derive the profit functions for each firm. The supplier's profits are:

$$\Pi_S = Q(w - c_S + (1 - \beta)y) - (1 - \alpha)f(y) - R,\tag{1}$$

the manufacturer's profits are:

$$\Pi_M = Q(p - w - c_M + \beta y) - \alpha f(y) + R,\tag{2}$$

and the supply-chain profits, equal to the sum of Π_S and Π_M, are:

$$\Pi_T = Q(p - c_S - c_M + y) - f(y).\tag{3}$$

These profit functions are relevant for examining both the non-cooperative and cooperative cases in the model.

The Non-Cooperative Case

In this case the firms do not contract on the supplier's R&D. The manufacturer therefore independently chooses α to maximize its own profits, and since R is paid from the supplier to the manufacturer, the supplier will set R equal to zero. The supplier then independently chooses y. A formal characterization of the solution is achieved by backward induction. The supplier's first-order condition, $\partial\Pi_S / \partial y = 0$, simplifies to:

$$f'(y) = \frac{Q(1 - \beta)}{1 - \alpha}.\tag{4}$$

Define $y_S(\alpha)$ as the solution to (4). The manufacturer's optimal α is then found by substituting $y_S(\alpha)$ into Π_M and solving the first-order condition $\partial\Pi_M / \partial\alpha = 0$. This simplifies to:

$$\frac{\partial y_S(\alpha)}{\partial\alpha}\left(Q\beta - \alpha\frac{\partial f(y_S(\alpha))}{\partial y}\right) - f(y_S(\alpha)) = 0 .^{9}\tag{5}$$

Define α_M as the solution to (5). The two-stage solution with no RJV is then characterized as $\alpha = \alpha_M$, $y = y_S(\alpha_M)$, and $R = 0$.

Some intuition helps illustrate the nature of this solution. If $\beta = 0$, none of the supplier's R&D benefit is passed along to the manufacturer via the component price. The manufacturer gets a zero return from its R&D investment, and we would expect α_M to equal zero. An increase in β raises the return to the manufacturer from the supplier's R&D, but it remains to be seen at what point α_M becomes positive. If $\beta = 1$ the manufacturer enjoys the entire cost benefit of the R&D, and α_M should

equal 1.

The Cooperative Case

Here the supplier and manufacturer form an RJV and jointly choose α and R to maximize combined profits and to satisfy the participation constraints, respectively. The supplier then independently chooses y. Define y_T as the efficient level of R&D; that is, the level of R&D that maximizes total supply chain profits. Note that y_T should be independent of β. The value β determines only how the gains from R&D are divided between the supplier and the manufacturer, and is therefore irrelevant when maximizing the sum of profits of the two firms. Next, define α_T as the fraction of R&D costs paid by the manufacturer such that the supplier chooses y_T. In characterizing these values, note first that $y_S(\alpha)$ remains unchanged from the non-cooperative case; that is, the supplier's R&D best response function to α is always the same. The first-order condition in the first stage then becomes $\partial\Pi_T/\partial\alpha = 0$, which simplifies to $\partial f(y_S(\alpha))/\partial y = Q$, and the solution is α_T.

Again, if $\beta = 0$ the supplier receives the entire benefit of the R&D, so α_T should equal zero. As β increases, so does the manufacturer's return from the supplier's R&D, and α_T increases to reflect this. When $\beta = 1$ the supplier realizes no benefit from the R&D, and α_T should equal 1. Proposition 1 summarizes this result.

Proposition 1. Define α_T as the value of α that induces the supplier to choose the efficient level of R&D. In equilibrium, $\alpha_T = \beta$.

Proof of Proposition 1. The supplier chooses the efficient level of R&D if its fraction of total R&D costs is equal to its fraction of the total benefits from R&D. The supplier's fraction of R&D costs equals $1-\alpha$. The total benefit of R&D is y per unit. The supplier nets only $(1-\beta)y$, however, because βy is passed on to the manufacturer through the reduced component price. The supplier's fraction of the total benefit therefore equals $1-\beta$, and it chooses the efficient level of R&D if $\alpha = \beta$.

Proposition 1 implies that an efficient outcome is achieved if the manufacturer pays a fraction of R&D costs that is equal to the fraction of the benefit it receives. However, while α_T ensures total profits are maximized, the value R must be set to ensure participation by both firms. Note first that there is a range of possible solutions. As long as R guarantees that each firm's profits are at least as high as under the non-cooperative case, the solution is acceptable and efficient. Because total profits must be higher in the cooperative case, R can be chosen to divide the increase in profits in such a way that both firms realize some of the gains. Proposition 2 summarizes this argument.

Proposition 2. If $\alpha = \alpha_T$ then the rebate makes the contract mutually agreeable to the supplier and the manufacturer if $R \in [\underline{R}, \overline{R}]$, where

$$\overline{R} = \Pi_S(\alpha_T, y_T) - \Pi_S(\alpha_M, y_S(\alpha_M)) \text{ and } \underline{R} = \Pi_M(\alpha_M, y_S(\alpha_M)) - \Pi_M(\alpha_T, y_T).$$

Proof of Proposition 2. The supplier is willing to pay the rebate R if $\Pi_S(\alpha_T, y_T, R) \geq \Pi_S(\alpha_M, y_S(\alpha_M))$. The highest the rebate can be is therefore $\overline{R}$, where $\Pi_S(\alpha_T, y_T, \overline{R}) \geq \Pi_S(\alpha_M, y_S(\alpha_M))$, or $\overline{R} = \Pi_S(\alpha_T, y_T) - \Pi_S(\alpha_M, y_S(\alpha_M))$. The manufacturer is willing to accept the rebate if $\Pi_M(\alpha_T, y_T, R) \geq \Pi_M(\alpha_M, y_S(\alpha_M))$. The lowest the rebate can be is therefore $\underline{R}$, where $\Pi_M(\alpha_T, y_T, \underline{R}) = \Pi_M(\alpha_M, y_S(\alpha_M))$ or $\underline{R} = \Pi_M(\alpha_M, y_S(\alpha_M)) - \Pi_M(\alpha_T, y_T)$.

Propositions 1 and 2 together show that if $\alpha = \alpha_T$ and R falls between $\underline{R}$ and $\overline{R}$, the contract induces the efficient level of R&D, and it is acceptable to both firms. Determining the actual value of R is a negotiation problem between the supplier and the manufacturer, so is not treated in this paper. However, note that the rebate is the sole determinant of the distribution of the profits gained from the contract. It is easy to see that $\overline{R} - \underline{R}$ equals the increase in total supply chain profits resulting from the contract. If $R = \underline{R}$, then the supplier obtains the entire gain. Increasing R shifts the gain from the supplier to the manufacturer, until $R = \overline{R}$, at which point the manufacturer obtains the entire gain. The realized value of R, and therefore the distribution of gains, depends on the relative bargaining power of the two firms.

Proposition 3. If $0 < \beta < 1$ then $\alpha_M < \alpha_T$. Also, if $\beta = 0$ then $\alpha_M = \alpha_T = 0$, and if $\beta = 1$ then $\alpha_M = \alpha_T = 1$.[10]

Proof of Proposition 3. The value α_M is determined by increasing α to the point at which the marginal net benefit to the manufacturer is zero. The value α_T is determined by increasing α to the point at which the marginal net benefit to the supply chain is zero. Because $y_S(\alpha)$ is monotonically increasing, it must be the case that α_M is less than α_T if the marginal net benefit to the manufacturer is less than that of the supply chain.

The benefit to the manufacturer is as follows: An increase in α induces an increase in y according to the function $y_S(\alpha)$. This provides a benefit to the manufacturer through the lower component price, $w - \beta y$. The benefit to the supply chain includes not only the manufacturer's benefit, but also that of the supplier. Firstly, as long as $\beta < 1$, an increase in α would benefit the supplier even if it did not change y. Of course the supplier increases y according to $y_S(\alpha)$, further improving its benefit. The benefit to the supply chain is therefore greater than that of the manufacturer alone. At α_M, when the net benefit to the manufacturer is zero, the net benefit to the supply chain must be positive. The value α_T must therefore be greater than α_M.

If $\beta = 0$, then $\alpha_T = 0$ because the manufacturer enjoys no benefit from the R&D. Mathematically, α_M would be negative. However, in practice the manufacturer could not enforce a negative α, so α_M would have to equal zero as well. If $\beta = 1$, the supplier does not realize any return from its R&D because all cost savings are passed along to the manufacturer. The supplier would therefore

choose $y = 0$ if it pays any of the R&D costs. The only way for the manufacturer to induce any R&D is to set α equal to 1, or in other words, $\alpha_M = 1$. In this case, because the supplier pays nothing for R&D and gets nothing in return, the net benefit to the manufacturer is equal to that of the supply chain. It is therefore the case that $\alpha_M = \alpha_T = 1$.

Proposition 3 demonstrates that if the manufacturer independently chooses α, then it usually pays too little of the supplier's R&D costs. Only at the extremes—when either none or all of the R&D benefit is passed along to the manufacturer—does the manufacturer choose the efficient fraction. These results have an obvious consequence on R&D and profits, as summarized by Proposition 4.

Proposition 4. If $0 < \beta < 1$ then $y_S(\alpha_M) < y_T$ and $\Pi_T(\alpha_M) < \Pi_T(\alpha_T)$. Also, if $\beta = 0$ or 1 then $y_S(\alpha_M) = y_T$ and $\Pi_T(\alpha_M) = \Pi_T(\alpha_T)$.

Proof of Proposition 4. These results follow directly from Proposition 3.

Proposition 4 demonstrates that cooperation has the greatest effect on R&D and profits when the amount of the R&D benefit that is passed along to the manufacturer is intermediate. The smallest effect occurs when either very little or very much of the benefit is passed along. In practice, little of the benefit might be passed along when the supplier faces little competition, when the R&D spillover rate is low, and/or when the manufacturer faces a high level of competition. The supplier may be forced to pass along most of the R&D benefit under the opposite conditions.

3. An Example

In this section an example is analyzed to illustrate the results. First, let us assume that $f(y) = ry^n$, with $n > 1$. The profit functions are then found by substituting this into (1), (2), and (3). From the supplier's first-order condition we find that $y_S(\alpha) = [Q(1-\beta)/nr(1-\alpha)]^{1/n-1}$. From the manufacturer's first-order condition it follows that:

$$\alpha_M = \frac{(2n-1)\beta - n + 1}{1 + (n-1)\beta}, \tag{6}$$

and then:

$$y_S(\alpha_M) = \left[\frac{Q(1 + (n-1)\beta)}{n^2 r} \right]^{\frac{1}{n-1}}. \tag{7}$$

A check shows that indeed $\alpha_T = \beta$ and

$$y_T = \left(\frac{Q}{nr} \right)^{\frac{1}{n-1}} .$$ (8)

To illustrate these results graphically, let us choose the following parameterization: $Q = 1000$, $p = 100$, $c_S = 20$, $c_M = 15$, $w = 50$, $n = 2$, and $r = 1$. The top graph in Figure 1 plots the equilibrium values α_M and α_T as β increases from 0 to 1. The middle graph plots the equilibrium values of $y(\alpha_M)$ and y_T, which are the levels of R&D under the non-cooperative and cooperative scenarios. The bottom graph depicts $\Pi_T(\alpha_M)$ and $\Pi_T(\alpha_T)$, which represent the supply-chain profits under the non-cooperative and cooperative scenarios.
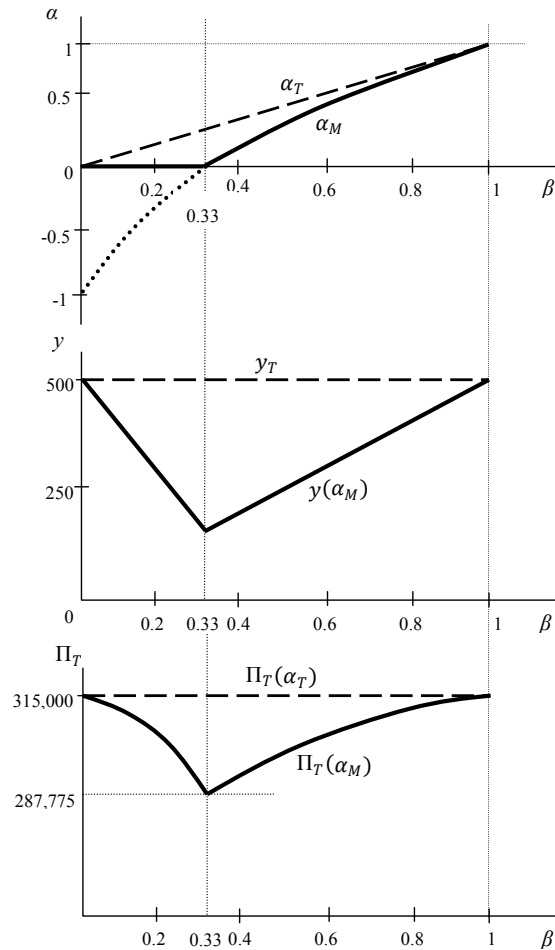
Figure 1. Equilibrium Values of α_M, α_T, $y(\alpha_M)$, y_T, $\Pi_T(\alpha_M)$, **and** $\Pi_T(\alpha_T)$

As predicted, the top graph shows that $\alpha_T = \beta$, and therefore α_T increases linearly from zero to one. Consistent with the proof of Proposition 3, the mathematical solution for α_M starts out negative at $\beta = 0$ and increases to equal one at $\beta = 1$. In this example, α_M would be negative for β between zero and about 0.33, and these values are illustrated with the heavy dotted segment of the α_M line. In reality, α_M would be zero over this range. The true solutions for α_M are thus illustrated by the thick black line, which equals zero for β up to 0.33, and then increases to one as β increases to one. Therefore, $\alpha_M = \alpha_T$ at the points $\beta = 0$ and $\beta = 1$, and $\alpha_M < \alpha_T$ for $0 < \beta < 1$. Importantly, the difference between α_M and α_T is greatest at the intermediate level $\beta = 0.33$, which is the point at which α_M becomes positive.

As the middle graph illustrates, $y(\alpha_M) = y_T$ when $\beta = 0$ and $\beta = 1$, and $y(\alpha_M) < y_T$ for $0 < \beta < 1$. Also, it can be seen that the gap between the non-cooperative and cooperative levels of R&D is highest at $\beta = 0.33$. The bottom graph shows a similar pattern with profits. There is no difference between the non-cooperative and cooperative cases when $\beta = 0$ and $\beta = 1$, and the largest difference again occurs at $\beta = 0.33$.

4. Conclusion

Given the legal flexibility afforded to vertically-related firms who wish to cooperate, the increased use of domestic and foreign outsourcing, and the relatively small amount of literature studying vertical research joint ventures, there seems to be a need for additional work in the area of research cooperation between buyers and their suppliers. To that end, the primary goal and innovation of this paper is to provide a theoretical basis and derivation of some conditions under which vertical research cooperation, in the form of R&D cost sharing and when vertical integration is not desirable, is beneficial.

The results of this paper indicate that if a supplier can undertake cost-reducing R&D, the buyer, or manufacturer, can benefit from a lower component price if it pays a fraction of the R&D costs. However, the manufacturer would choose a fraction that is too little to induce the supplier to undertake the level of R&D that maximizes supply chain profits. If the manufacturer increases this fraction to equal the percentage of the R&D benefit that is passed on to the manufacturer through the component price, the supplier would choose the optimal level of R&D. To ensure individual rationality, the supplier can pay a fixed fee to the manufacturer that is independent of the amount of research it conducts.

While cooperatively choosing the manufacturer's share of R&D costs (almost) always improves total profits, the magnitude of the effect depends on how much of the R&D benefit is passed on to the manufacturer. Cooperation has the greatest effect when the fraction passed on is intermediate—not extremely low or high. Such low or high fractions may be caused by a number of factors, including R&D spillover rates and the levels of competition in the supplier's and manufacturer's markets. For example, if the supplier operates in a highly competitive market and if

the R&D spillover rate is high, we would expect a large fraction of the cost reduction to be passed on. In this case, R&D cooperation would not have a great impact on overall profits.

Because the research on R&D cooperation between vertically-related firms is not yet fully developed, there are many avenues for future research. For example, this paper focused on cost-reducing R&D, but cooperation in the design phase or specifications of components may be as or more important under various circumstances. There are also intellectual property and patent implications, the magnitude of which may depend also on the competitive environments, spillover rates, international comity and patent law, and other factors. As such, the present paper can hopefully lay some foundation for future research that attends to these issues.

Notes

1. See, for example, Gray et al. (2007), who support this claim with an empirical analysis of the product quality differences between a vertically integrated supplier and an independent supplier.

2. In the US and many other countries, vertically-related firms are permitted even to coordinate prices and output, as demonstrated by the common use of vendor-managed inventory systems, resale price maintenance, and price coordination contracts. Restrictions on joint ventures between vertically-related firms are limited mostly to their effects on the cost of entry into upstream or downstream markets.

3. Ideally, coordination of actions between vertically-related firms is best achieved by integration. However, as described by Williamson (1971, 1991), high agency efficiency relative to technical complementarity between vertically-related firms can make integration inefficient. Grossman and Hart (1986) add that vertical integration is less attractive when the impacts of the two firms' relationship-specific investments are comparable. This is assumed to be the case in the present paper and in the related research.

4. There is a related body of literature on the effects of sharing a firm's productivity and cost information with its suppliers and/or its competitors. See Creane (2007, 2008), Zhang (2002), and Li (2002) for recent studies on this topic. Also, the "holdup problem" is not applicable here, since the manufacturer is not permitted to renegotiate the component price after the supplier makes its R&D investment.

5. There may be other suppliers willing to produce the component at a lower price should they easily learn the improved production methods. The current supplier may need to lower its price to some degree to prevent this. Also, a manufacturer facing little competition would have more power in negotiating a component price.

6. Of course the manufacturer may also need to pass some of its cost savings on through its output price, but this assumption is not necessary to illustrate the effects of an RJV on R&D. Indeed, lowering the manufacturer's output price somewhat would add a benefit to consumers, but it does not change the nature of the results.

7. It is assumed that y itself is not verifiable in courts and thus cannot be directly contracted on. On the other hand, R&D expenditures can be contracted on because they can be made observable to both firms, and in fact may require financial commitment even before the R&D takes place.

8. The value w is the component price with no R&D. The value β represents the fraction of the supplier's cost reduction that is passed along to the manufacturer via the component price. As stated, a higher β can be the result of more competition in the supplier's market, less competition in the manufacturer's market, or a higher R&D spillover rate. Again, the "holdup problem" is not an issue here because β is determined exogenously.

9. A check shows that conditions (4) and (5) represent maxima. The term $\partial^2 \Pi_S / \partial^2 y^2 = -f''(y)(1-\alpha)$ is negative, showing that (4) represents a maximum. Also, substituting y_s into the supplier's profit function yields $\Pi_S = Q(w - c_s) + Q(1 - \beta)y_s(\alpha) - (1 - \alpha)f(y_s(\alpha)) - R$. Because $f'(y) > 0$ and $f''(y) > 0$, the increase in the rate of increase of $f(y_s(\alpha))$ is greater than that of $y_s(\alpha)$. This demonstrates that Π_S is concave and that (5) represents a maximum. A similar proof shows that the first-order conditions in the cooperative case represent maxima as well.

10. An irregularity is that if $\beta = 1$ and the manufacturer pays for all of the R&D, the supplier has no incentive to choose y_T. The manufacturer could solve this problem by, for example, setting α slightly more than one for R&D up to y_T, and zero after that. Also, if $\beta = 0$, the mathematical solution for α_M would be negative, meaning that in practice α_M would also be zero.

References

Banerjee, S. and P. Lin, (2001), "Vertical Research Joint Ventures," *International Journal of Industrial Organization*, 19, 285-302.

Creane, A., (2007), "Productivity Information in Vertical Sharing Agreements," *International Journal of Industrial Organization*, 25, 821-841.

Creane, A., (2008), "Input Suppliers, Differential Pricing, and Information Sharing Agreements," *Journal of Economics & Management Strategy*, 17, 865-893.

d'Aspremont, C. and A. Jacquemin, (1988), "Cooperative and Noncooperative R&D in Duopoly with Spillovers," *American Economic Review*, 78, 1133-1137.

Gray, J., A. Roth, and B. Tomlin, (2007) "Quality Risk in Outsourcing: Evidence from the U.S. Drug Industry," mimeo.

Grossman, S. and O. Hart, (1986), "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration," *Journal of Political Economy*, 94, 691-719.

Ishii, A., (2004), "Cooperative R&D between Vertically Related Firms with Spillovers," *International Journal of Industrial Organization*, 22, 1213-1235.

Kamien, M., E. Muller, and I. Zang, (1992), "Research Joint Ventures and R&D Cartels," *American Economic Review*, 82, 1293-1306.

Li, L., (2002), "Information Sharing in a Supply Chain with Horizontal Competition," *Management Science*, 48, 1196-1212.

Williamson, O., (1971), "The Vertical Integration of Production: Market Failure Considerations," *American Economic Review*, 61, 112-123.

Williamson, O., (1991), "Strategizing, Economizing and Economic Organization," *Strategic Management Journal*, 12, 75-94.

Zhang, H., (2002), "Vertical Information Exchange in a Supply Chain with Duopoly Retailers," *Production and Operations Management*, 11, 531-546.