Kogut, B. and S.J. Chang, 1991, "Technological Capabilities and Japanese Direct Investment in the United States," *Review of Economics and Statistics*, 73, 401–413.
—— and ——, 1996, "Platform Investments and Volatile Exchange Rates: Direct Investment in the U.S. by Japanese Electronic Companies," *Review of Economics and Statistics*, 78, 221–231.
Motta, M., 1994, "International Trade and Investments in a Vertically Differentiated Industry," *International Journal of Industrial Organization*, 12, 179–196.
Ottaviano, G., T. Tabuchi, and J.-F. Thisse, 2002, "Agglomeration and Trade Revisited," *International Economic Review*, 43 (2), 409–436.
Varian, H., 1992, *Microeconomic Analysis*, New York: W.W. Norton.
Yu, C.-M.J. and K. Ito, 1988, "Oligopolistic Reaction and Foreign Direct Investment: The Case of the U.S. Tire and Textiles Industry," *Journal of International Business Studies*, 19, 449–460.

Why Should the Boss Own the Assets?

Birger Wernerfelt
Sloan School
Massachusetts Institute of Technology
Cambridge, MA 02142
bwerner@mit.edu

In the context of an employment relationship, I present an argument suggesting that it is more efficient for the boss to own the productive assets. The idea is that a conflict between productivity and depreciation is internalized if the player deciding what an asset is used for also has residual claims. An empirical test finds evidence consistent with this. By asking whether the boss should own the assets, the paper reverses the reasoning from the literature in which it is argued that the owner has power and thus is the boss.

1. Introduction

In employment relationships, the owner of the company most often owns the productive assets used by employees (Holmstrom, 1999).[1] There are some exceptions to this, mostly in cases where only one operator uses the asset (examples include chefs owning knives, hairdressers owning scissors, auto mechanics owning tools, and salespeople owning cars). However, these exceptions are so few that the Internal Revenue Service (U.S. Department of the Treasury, 1990) lists nonownership of assets as one of twenty diagnostic criteria for identifying employee status.

Grossman and Hart (1986) and Hart and Moore (1990) offer one explanation for this stylized fact. They show that ownership confers bargaining power and therefore furthers investment incentives under incomplete contracting. In a logically independent, but consistent, step, they then define boss as the player owning the asset. A broader story is told by Holmstrom and Milgrom (1994), who look at

I am grateful for comments by two anonymous referees, Oliver Hart, Duncan Simester, and Dan Spulber, as well as seminar participants at Duke, Harvard, Iowa, MIT, and Washington University. The usual disclaimer applies.

1. In many examples the manager is the agent of the owners, but we will follow the literature and defer this complication for later study.

Journal of Economics & Management Strategy, Volume 11, Number 3, Fall 2002, 473–485

asset ownership as an element of incentive systems for multitasking agents. They argue that several dimensions of these incentive systems can be expected to covary, and define an employee as an agent who is subject to a "firm" incentive system. This system is again defined by several characteristics, but one of them is that the employer owns the assets.

While these explanations generally are perceived as somewhat different, they have a very similar perspective on the function of asset ownership. In fact, Holmstrom and Milgrom (1994, p. 975) cite Grossman, Hart, and Moore in arguing that ownership confers incentives because it gives *ex post* bargaining power. So the idea in the literature is that ownership gives power and that this power, among other things (Holmstrom and Milgrom, 1994, p. 989), justifies affixing the boss label to the owner.

In the present paper, I examine the linkage in reverse. Assuming that a player has the right to give orders about asset use, I show that this player is also likely to have residual claims on (own) the assets. The argument revolves around the externality between the productivity and depreciation of the asset. The idea is that if the manager also has residual claims on an asset and thus bears the costs of depreciation, she can take them into account when deciding how the asset is used.

I assume that the parties can contract on measures of depreciation and show that ownership can be a substitute for contracts, such that a player whose actions have a less contractible effect on depreciation is more likely to own the asset. Similarly, I show that players who have a greater effect on depreciation are more likely to own the asset. This implies that the manager is more likely to have residual claims on assets that are used by several employees. The theory depends on unforeseen contingencies, but has the appealing property that residual claims have a role even if the parties can contract on dollar payoffs.[2]

The model is presented and analyzed in Section 2, and some empirical evidence is offered in Section 3. The results are summarized and related to the theory of the firm in Section 4.

2. The Externality Argument

This section develops the model, concentrating on an example with two players and a single asset. A manager requires a service from a physical asset, but needs an employee to operate the asset. The players must agree on a contract before the manager tells the employee what specific service she wants.[3] The employee can use the asset to produce a single service drawn from a very large set A. The set A is known *ex ante*, but only some of its elements will turn out to be feasible. Associated with each feasible service is the level of activity, $a \in R$, that it requires. Conversely, for any real a, it is known that there will be a feasible service corresponding to that a. To keep the analysis simple, I assume that all payoff-relevant implications of any specific service can be summarized by two functions of a. The benefits, which accrue to the manager, are given by $B(a)$, and the depreciation, which is borne by the owner of the asset, is given by $D(a)$. While the players do not *ex ante* know the set of feasible services, they know that there will be frontier along which the monetary equivalent of the benefit is $B(a)$, while the depreciation costs are $D(a)$. After contracting, the manager learns the identity of the feasible set of services, while the employee only learns the identity of the specific element, indexed by a^*, that he is asked to provide.

Assume that the $B(a)$, $D(a)$ frontier is stochastic, so that the benefits of incurring higher depreciation costs are *ex ante* uncertain. To keep the exposition as transparent as possible, I will use very simple functional forms for $B(a)$ and $D(a)$. Specifically, the benefits are assumed to be a realization of

$$B = (1 + \varepsilon_b)a - \tfrac{1}{2}a^2, \qquad \varepsilon_b \sim N(0, \sigma_b^2), \tag{1}$$

and the depreciation is a realization of

$$D = \phi_a a - \phi_m m + \varepsilon_d, \qquad \varepsilon_d \sim N(0, \sigma_d^2), \quad \phi_a, \phi_m \in R_+^2, \tag{2}$$

where m is the level of effort the employee puts into maintenance. The parameters ϕ_a and ϕ_m reflect the extent to which depreciation is sensitive to the level of activity chosen by the boss and the maintenance effort of the employee.

The realization of the benefit noise ε_b is the most important uncertainty, affecting the relative efficiency of different values of a. In particular, when ε_b is higher, it is efficient to incur higher depreciation costs. I assume that the set of feasible services reveals ε_b and ε_d to the boss, and that any element of this set reveals a and ε_d to the employee. The trade-off between depreciation and benefits is illustrated in Figure 1, where the arrows indicate the effects of different realizations of ε_d and ε_b, taking into account that a large ε_b makes a larger a and thus a larger D attractive.

2. It has been argued that other theories of ownership are vulnerable to the ability to write such contracts (Maskin and Tirole, 1999a,b; Hart and Moore, 1999; Wernerfelt, 1989).

3. This arrangement is consistent with the model of Wernerfelt (1997) in which an employment relationship is well suited to adaptation under uncertainty.

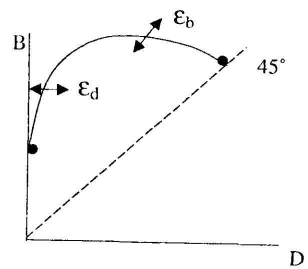

FIGURE 1. THE BENEFIT-DEPRECIATION FRONTIER

To make clear that the argument does not depend on multitasking, I assume that the employee's effort costs C depend only on maintenance efforts and not on the level of productive activities, a:

$$C = \tfrac{1}{2}m^2. \tag{3}$$

The employee may still care about a, because he may bear depreciation costs, but as he does not know the set of feasible services himself, he has to learn about them from the boss.

Assume that the players may contract on the following verifiable measures of activity level and efforts:

$$x_a = a + \varepsilon_a, \qquad \varepsilon_a \sim N(0, \sigma_a^2) \tag{4}$$

$$x_m = m + \varepsilon_m, \qquad \varepsilon_m \sim N(0, \sigma_m^2). \tag{5}$$

The contractibility of x_a, together with the construction of $B(a)$ and $D(a)$, gives the players the ability to contract on (dollar) payoffs in spite of the unforeseen contingencies. The results do not require that $\sigma_a, \sigma_m > 0$, so payoffs may be perfectly contractible.

I restrict attention to linear contracts, such that the employee's gross pay is the sum of

$$s_a = \alpha_a x_a + \beta \tag{6}$$

and

$$s_m = \alpha_m x_m, \tag{7}$$

where α_a, α_m, and β are constants.

Still following the literature, I assume that the manager is risk-neutral, while the employees value payments v as $-\exp(-rv)$, $r > 0$. Under known, but tortured, assumptions, the linear contract (6) and (7) may then be optimal in a dynamic version of the model

(Holmstrom and Milgrom, 1987). The linearity will help keep the analysis transparent, but the intuition behind the results does not depend on this feature.

In this model the owner of the asset is the residual claimant, the player who bears the depreciation costs D. I use the parameter $\lambda \in [0, 1]$ to indicate the degree to which the employee owns the asset. To keep the analysis simpler, I treat ownership as a continuous variable. For reasons outside the model, it may be reasonable to think of λ as a latent variable determining discrete ownership for some critical value. However, the continuous results are interesting, and, much as in the franchising literature (Bhattacharyya and Lafontaine, 1995), the results will suggest that an intermediate value may be the best way to balance incentives.

Using the notation of the model, the sequence of events and the information structure is as follows:

0. The players agree on a contract $(\alpha_a, \alpha_m, \beta)$ and allocate asset ownership (λ).
1. The manager learns the set of feasible services and the realizations of ε_b, ε_d.
2. The manager requests a specific service indexed by a^*.
3. The employee learns a^* and ε_d, and selects m.
4. The random variables ε_a and ε_m are realized.
5. Payoffs s_a, s_m, λD, $(1-\lambda)D$, B, and C are made.

This extensive form assumes that the original contract is binding for both players. Neither can opt out, and there is no possibility of renegotiation or reallocation of ownership. Given this, the employee's problem is to maximize the expected utility of

$$\alpha_a x_a + \beta + \alpha_m x_m - \tfrac{1}{2}m^2 - \lambda(\phi_a a - \phi_m m + \varepsilon_d). \tag{8}$$

For given a, ε_b, ε_d, and independent ε's, standard results give us the *ex interim* certainty equivalent of (8) as

$$\alpha_m m + \beta + \alpha_a a - \tfrac{1}{2}m^2 - \lambda(\phi_a a - \phi_m m - \varepsilon_d) - \tfrac{1}{2}r\left(\alpha_m^2\sigma_m^2 + \alpha_a^2\sigma_a^2\right). \tag{9}$$

Therefore, by differentiation,

$$m^* = \alpha_m + \phi_m\lambda. \tag{10}$$

Given this behavior, the manager wants to maximize the expectation of

$$-\alpha_a a - \beta - \alpha_m(\alpha_m + \phi_m\lambda) - (1-\lambda)(\phi_a a - \phi_m\alpha_m - \phi_m^2\lambda + \varepsilon_d) + (1+\varepsilon_b)a - \tfrac{1}{2}a^2. \tag{11}$$

Since a is chosen after ε_b is known, the first-order condition yields the index of the desired service as

$$a^* = \lambda\phi_a - \phi_a - \alpha_a + 1 + \varepsilon_b. \tag{12}$$

The first-best level of activity is $1 + \varepsilon_b - \phi_a$, so it is distorted upward by employee ownership (λ) and downward by the compensation (α_a).

The *ex ante* certainty equivalent of total surplus, taking into account that the employee bears part of the risk associated with different levels of depreciation (generated by ε_b and ε_d), is given by

$$\begin{aligned}&\lambda\phi_a - \phi_a - \alpha_a + 1 + \sigma_b^2 - \tfrac{1}{2}(\lambda\phi_a - \phi_a - \alpha_a + 1)^2 - \tfrac{1}{2}\sigma_b^2 \\ &- \lambda\phi_a^2 + \phi_a^2 + \phi_a\alpha_a - \phi_a + \phi_m\alpha_m + \phi_m^2\lambda - \tfrac{1}{2}(\alpha_m + \phi_m\lambda)^2 \\ &- \tfrac{1}{2} r\left[\alpha_m^2\sigma_m^2 + \alpha_a^2(\sigma_a^2 + \sigma_b^2) + \lambda^2(\phi_a^2\sigma_b^2 + \sigma_d^2)\right].\end{aligned} \tag{13}$$

The optimal contract α_a^*, α_m^* and the optimal ownership allocation λ^* can be found by maximization of (13). The first-order conditions yield

$$\alpha_a^* = \frac{\lambda^*\phi_a}{1 + r(\sigma_a^2 + \sigma_b^2)}, \tag{14}$$

$$\alpha_m^* = \frac{(1-\lambda)\phi_m}{1 + r\sigma_m^2}, \tag{15}$$

$$\lambda^* = \frac{\phi_m^2 + \phi_a\alpha_a^* - \phi_m\alpha_m^*}{\phi_m^2 + \phi_a^2 + r(\phi_a^2\sigma_b^2 + \sigma_d^2)}. \tag{16}$$

Trivial, but tedious substitution now yields

PROPOSITION 1: *The employee ownership ratio λ^* takes values in the unit interval and is decreasing in the benefit variance σ_b; decreasing in the measurement variance on the activity level chosen by the manager, σ_a; increasing in the measurement variance on the employee's maintenance effort, σ_m; decreasing in the effect of the activity level on depreciation, ϕ_a; increasing in the effect of maintenance effort on depreciation, ϕ_m; and decreasing in the depreciation variance σ_d.*

The central result in the paper is that the employee ownership ratio is lower when the benefit variance is larger. This means that the externality between productivity and depreciation is internalized if the manager owns the assets. Since the benefits of putting extra wear and tear on the asset vary and only she knows the trade-off, the team is going to implement better adjustments if she owns. The alternative is to compensate her by contracting on a. However, this cannot work perfectly, because the variation in the marginal returns to a, viz. ε_b, is a risk from the perspective of the employee. Consistent with this, we see from (12) and (14) that a^* equals the first-best level when $\lambda = 0$. The reason that $\lambda^* > 0$ is that ownership also influences the maintenance incentives of the employee. When these can be handled contractually, or do not need to be handled, the model suggests that the boss should own the asset outright. That is, $\lambda^* = 0$ if $\sigma_m = 0$ or $\phi_m = 0$.

The effects of errors in measurements of the manager's chosen level of activity and the employee's maintenance effort are parallel. They say that if the depreciation resulting from a player's decisions becomes more contractible, the player is less likely to own the asset. Ownership is a substitute for a contract. As an example, if my manager tells me how to use my car, she will most likely compensate me on a per-mile basis. To the extent that miles driven are a poor measure of depreciation caused by different uses of the car, it will be more efficient for the manager to own the car. In the limit when the employee's maintenance efforts are perfectly contractible, there is no need for him to own the asset, and $\lambda^* = 0$. If the manager's activity selections are perfectly contractible, $\lambda^* < 1$ because the employee still has to be compensated for the risk brought on by ε_b.

The effects of the sensitivity of depreciation to different activity levels and maintenance efforts are also parallel. They say that if a player's decisions affect depreciation more, then that player is more likely to own the asset. This effect is mentioned by Hart (1995, p. 50), who says that the asset should be owned by the person who has the larger influence on its depreciation. For example, ownership by the manager is favored if the decisions she is likely to make have steep implications for asset depreciation. At the moment, a violinist in the Boston Symphony Orchestra owns his violin. The manager (director) will ask the violinist to use the instrument in many different ways, but there is little wear and tear ($\phi_a = 0$). In fact, the violinist is not even compensated for any measure of depreciation, like x_a in the model. Suppose on the other hand, that the director borrowed a page from Jimi Hendrix and occasionally asked the violinist to set fire to the instrument. In such a case the parties would have to do one of two things: either find a good contractible measure of depreciation (with low σ_a) or have the boss own at least part of the asset. Conversely, if the employee has no influence on depreciation ($\phi_m = 0$), then the boss should own the asset.

The effect of depreciation variance simply says that the less risk-averse player is better able to absorb depreciation risks associated with ownership. Ownership plays a subtle role in the model. As is well known, contracting costs appear as risk costs in a model of this type. The players' influences on depreciation, a and m, are imperfectly contractible, and the owner bears that part of depreciation that is not captured by the contractual payments, s_a and s_m. Because the players correctly anticipate the distribution of depreciation costs, the payments are correct on an expected-value basis. For example, if the employee owns the asset, the manager has incentives to select higher values of a, thereby causing the asset to depreciate faster. The optimal contracts anticipate such mean effects of the externality, compensating for a higher λ by a higher α_a, and thus paying the employee more for higher levels of a. The noncontractible depreciation costs are the (risk) costs of the error variance in the contracts compensating owners for depreciation. The manager could help an employee-owner by reducing the variance in depreciation, but this externality is not captured, because of contracting limitations. To mitigate this problem, more of the residual claims are allocated to the manager.

It is critical that one of the players be risk-averse. With two risk-neutral players, even the linear contract (6) and (7) can implement the first best for any allocation of ownership. I have followed convention and assumed that the employee is risk-averse, while the boss is risk-neutral. Since the boss often is a corporation or a wealthier individual, this seems a reasonable first cut. If the boss is risk-averse, the results are symmetric. The optimal value of $1 - \lambda$ is given by the right side of (16) with the subscripts a and m interchanged.

For an asset used by a single employee, the results indicate that some forces favor ownership by the manager and others favor ownership by the employee. Simester and Wernerfelt (2001) suggest and empirically show that the balance of these forces changes as the number of operators goes up. To look at this in the context of the model, I can scale B and D by the number of employees, n, to get

$$B = n(1 + \varepsilon_b)a - \frac{n}{2}a^2, \qquad \varepsilon_b \sim N(0, n\sigma_b^2) \tag{17}$$

and

$$D = n\phi_a a - \phi_m \sum m + \varepsilon_d,\ \varepsilon_d \sim N(0, n\sigma_d^2), \qquad \phi_a, \phi_m \in R_+^2. \tag{18}$$

In this version of the model, the share of assets owned by the employees should go to zero as n goes to infinity. Intuitively, the importance of individual employees goes down, while the relative importance of the boss remains the same. So assets that are used by several employees are more likely to be owned by the manager. In support of this result, it is interesting that almost all examples of employee-owned assets are used exclusively by their owner.

3. Evidence

I now provide some empirical support for the argument. As stated in the premise of the paper and elsewhere, there are few examples of productive assets owned by employees. It is even harder to find examples of assets that only sometimes are owned by employees (and otherwise by employers). Automobiles used by salespeople are, however, a very good example. In a study of 457 US companies, Watson Wyatt (1997) found that sales employees owned or leased their cars in 40% of the cases, while the firms owned or leased them in 44% of the cases (16% used a combination). I will now try to see if the theory from Section 2 can explain interfirm differences in this area. The hypothesis is that firms are more likely to own the cars if they need to give their salespeople more directions about what to use them for.

I am not aware of any other studies of this particular issue, and it is not clear what other theories of asset ownership would predict. However, concerns about risk sharing or tax advantages would seem to favor ownership by the firm, independent of the extent to which it gives directions to the employee. In more interesting contrast to the theory in the present paper, believers in the Grossman-Hart-Moore theory would presumably predict that firms that own the cars are more likely to give directions about how they are used (so once again the direction of causality is reversed).

Textbooks in sales management rarely discuss this, and when they do, the main point seems to be that firm ownership is "better," but only sufficiently large firms should do it (Stanton and Buskirk, 1987). The reasons given are not perfectly clear but appear to be related to administrative economies of scale and capacity utilization. From rent-versus-buy decisions on vacation homes and cars abroad, it seems clear that some fixed costs are involved in ownership. However, I will not pursue that aspect here.

The determination of ownership of the cars used by salespeople is a very good fit with the model from Section 2. The depreciation of the cars depends on how well the drivers take care of them and how the work requires them to be driven. Furthermore, most salespeople are paid on a linear mileage schedule to compensate for wear and tear on cars owned by them (Runzheimer, 1996), but the milage compensation is noisy (as witnessed by the fact that the prices of used cars are imperfectly predicted by their odometer readings).

By looking at a single class of assets, we hold constant a number of factors bearing on ownership. Ideally, I would like to interpret any variation in ownership as a result of interfirm differences in the extent to which the firm exercises significant control over how the car is used and thus how much it depreciates.

To pick up this variation, we conducted a survey of managers identified in the *Directory of Corporate Affiliations* (1997). I randomly selected 146 firms in *Volume IV: US Private Companies*, and sent a personalized letter to a manager, if possible a sales manager, in each. A cover letter stated that we were interested in identifying best practices and promised anonymity. The questionnaire tries to measure the extent to which the firm gets involved in the salesperson's use of the car, by asking whether he is instructed to visit specific customers. The precise wording is given in Table I, and the hypothesis is, of course, that firms that often make such instructions are more likely to own the cars used by their salespeople. (In a smaller pilot study, I also asked whether the salespeople made their own call plans. However, I did not get any variance on that item. It seems that almost all salespeople do this detailed scheduling themselves.)

I got 45 responses, for a response rate of 31%. Given that several firms could not reply because they used independent sales reps, the response rate is good for studies of this type [Wernerfelt (1997) got 25%]. On item **a**, the mean response is 5.8, with a standard deviation of 1.6, while on item **b** it was 68, with a standard deviation of 46. As the rating on item **a** goes up by one point, an additional 5% of the cars were, on the average, company-owned. Because of the nature of the scales, we need to use the rank correlation between the two items to test the hypothesis. Kendal's τ_b is 0.279, significant at the .02 level (one-tailed test). So although it is weak, this statistical evidence is consistent with the theory.

Since I only have a correlation, it is possible to read the causality as going the other way. That is, following Grossman, Hart, and Moore, one could suggest that the data show that firms, once they provide

TABLE I.
QUESTIONNAIRE

For those members of your sales force who are direct employees (as opposed to independent reps) please mark the number below that best describes your opinion.

a. We often instruct individual salespeople to visit specific customers.

strongly disagree			neutral			strongly agree
1	2	3	4	5	6	7

b. What percentage of them use a car which is owned or leased by the company (as opposed to owned or leased by the employee)? ________%

company cars, are more aggressive about giving instructions. The test does not discriminate between their theory and that presented here; it fails to reject either. On the other hand, as soon as employer status is determined by factors beyond asset ownership, one can imagine a model with simultaneous allocation of employer-employee roles and asset ownership. In such a model the two would then be complements. So even if the main determinant of ownership is not that argued here, the data suggest that the manager's right to give instructions is relevant to the allocation of ownership.

4. Discussion

This discussion contributes to the study of ownership of productive assets. When contracts are incomplete, the literature has identified many functions of ownership. Holmstrom and Milgrom (1994, p. 972) write that it helps give operators incentives to maintain the assets they work with, Holmstrom (1999) has argued that it allows the boss to design better incentive systems, and Grossman and Hart (1986) and Hart and Moore (1990) have shown that it helps implement specific human-capital investments. I have here added another function, which becomes a logical possibility once you break the definitional link between asset ownership and the role of boss in the employment relationship. Specifically, I show that asset ownership may help give proper incentives to the agent controlling the use of an asset (the manager or the boss). There is, however, no reason to believe that our understanding of the functions of ownership is complete. Much work remains in expanding the list of factors bearing on ownership and empirically evaluating the relative importance of the individual items on it.

The theory has interesting implications for the definition of a firm. In the Grossman-Hart-Moore model, the firm and its scope are defined by a set of co-owned assets. In Holmstrom and Milgrom (1994) and Wernerfelt (1997), it is defined by a set of agents who sell their human asset services to a single boss under a specific contract. Holmstrom and Milgrom define the contract as one in which, *inter alia*, the employee does not own the assets, while Wernerfelt focuses on the agreement to follow *ex ante* unspecified orders. The present paper complements the latter definition of the firm by explaining the empirical covariation between the role of boss and asset ownership.

Corresponding to the different definitions of the firm are different views of the boss. In the Grossman-Hart-Moore and Holmstrom-Milgrom models, the boss has power (Rajan and Zingales, 1998), and she is selected for the job because her incentives are important. In contrast, the model in the present paper is consistent with the more classical view (Simon, 1951) that the role of boss reflects superior information [much like "real authority" in Aghion and Tirole (1997)]. That is, the employment relationship is not primarily about allocating power to give incentives, but about letting the best-informed player decide.

It is possible to interpret my externality argument as saying that the manager is making a specific investment if she chooses to use the asset in ways that entail fewer depreciation costs. With this interpretation, the theory in the present paper almost exactly reverses the causality of the Grossman-Hart-Moore theory. Specifically, the argument is that because she is the manager, a player is in a position to make the specific investment and that this then makes it more efficient for her to own. Grossman, Hart, and Moore start with a technologically determined need to make a specific investment, and derive the ownership, and the organizational role as manager, from that.

On a more technical level, a number of papers (Maskin and Tirole, 1999a,b; Tirole, 1999; Wernerfelt, 1989) have taken issue with the modeling of incomplete contracts in general and unforeseen contingencies in particular (see also Hart and Moore, 1999). The idea is that unforeseen contingencies do not prevent the writing of contracts on payoffs, since these take values in a known space (the real line). This has cast doubt on the justification for incomplete contracts and the foundation for any theory of ownership. In the present paper, I do allow the players to write contracts on the dollar payoffs from *ex ante* unforeseen services. However, I do not allow them to contract on the variance in these payoffs (σ_b in the model). This inefficiency then gives a role for ownership. While it is an unpleasant feature of this modeling technique that ownership becomes irrelevant if neither player is risk-averse, the paper shares this drawback with much of the literature.

References

Aghion, P. and J. Tirole, 1997, "Formal and Real Authority in Organizations," *Journal of Political Economy*, 105 (1), 1–29.

Bhattacharyya, S. and F. Lafontaine, 1995, "Double-Sided Moral Hazard and the Nature of Share Contracts," *RAND Journal of Economics*, 26 (4), 761–781.

Directory of Corporate Affiliations IV: US Private Companies, 1997, New Providence, RI: National Register Publications.

Grossman, S. and O.D. Hart, 1986, "The Costs and Benefits of Ownership," *Journal of Political Economy*, 97 (4), 691–719.

Hart, O.D., 1995, *Firms, Contracts, and Financial Structure*, Oxford, UK: Oxford University Press.

——, and J. Moore, 1990, "Property Rights and Nature of the Firm," *Journal of Political Economy*, 98 (6), 1119–1158.

—— and ——, 1999, "Foundations of Incomplete Contracts," *Review of Economic Studies*, 66 (1), 115–138.

Holmstrom, B., 1999, "The Firm as a Subeconomy," *Journal of Law, Economics, and Organization*, 15 (1), 74–102.

—— and P. Milgrom, 1987, "Aggregation and Linearity in the Provision of Intertemporal Incentives," *Econometrica*, 55 (2), 303–328.

—— and ——, 1994, "The Firm as an Incentive System," *American Economic Review*, 84 (4), 972–991.

Maskin, E. and J. Tirole, 1999a, "Unforeseen Contingencies and Incomplete Contracts," *Review of Economic Studies*, 66 (1), 83–114.

—— and ——, 1999b, "Two Remarks on the Property-Rights Literature," *Review of Economic Studies*, 66 (1), 139–150.

Rajan, R.G. and L. Zingales, 1998, "Power in a Theory of the Firm," *Quarterly Journal of Economics*, 113 (2), 387–432.

Runzheimer International, 1996, *Financial Treatment of Executive and Business Vehicles*, Runzheimer Park, WI.

Simester, D.I. and B. Wernerfelt, 2001, "Determinants of Asset Ownership: A Study of the Carpentry Trade," Mimeo, MIT.

Simon, H.A., 1951, "A Formal Theory of the Employment Relationship," *Econometrica*, 19 (4), 293–305.

Stanton, W.J. and R.H. Buskirk, 1987, *Management of the Salesforce*, Homewood, IL: Irwin.

Tirole, J., 1999, "Incomplete Contracts: Where Do We Stand?" *Econometrica*, 67 (4), 741–782.

U.S. Department of the Treasury, Internal Revenue Service, 1990, *Business Reporting Publication 937*, Washington: U.S. Department of the Treasury.

Watson Wyatt Data Services, 1997, *The 1997/98 ECS Survey of Sales and Marketing Personnel Compensation*, Rochelle Park, NJ: Watson Wyatt.

Wernerfelt, B., 1989, "Unforeseen Contingencies and Market Failure," Mimeo, MIT.

——, 1997, "On the Nature and Scope of the Firm: An Adjustment-Cost Theory," *Journal of Business*, 70 (4), 489–514.