Resolving Double Moral Hazard Problems with Buyout Agreements

by

Joel S. Demski*

and

David E. M. Sappington**

Revised March 1991

* Yale University.
** University of Florida.

We are grateful to Debra Aron, Peter Cramton, Jerry Feltham, Robert Gibbons, Jon Ingersoll, Michael Riordan, Richard Romano, Steve Ryan, Alan Schwartz, and particularly Benjamin Hermalin for very helpful comments. Financial support from the National Science Foundation and the Public Utilities Research Center at the University of Florida is gratefully acknowledged.

ABSTRACT

We consider a double moral hazard problem in which the efforts of two parties, e.g., a principal who initially owns an enterprise and a risk averse agent in the enterprise, are not verifiable. The realized value of the enterprise's random profit stream is also unverifiable. There is also no third party to break a "balanced budget" requirement. Nevertheless, the double moral hazard problem can be resolved completely and costlessly when the principal, who can observe the agent's actions, has the option of requiring the agent to purchase the enterprise at a prenegotiated price.

1. Introduction.

It has long been recognized that workers whose wages are not linked to the firm's performance may have limited incentive to work diligently, particularly when observed shirking by an employee cannot be documented conclusively to any third party (i.e., when shirking is not verifiable). Of course, some incentive for workers to labor diligently may be restored if they "buy out" their employer, and become the residual claimants for the firm's stream of profits.

In this article, as one application of a more general phenomenon, we propose an expanded role for labor buyouts. We document the gains that arise when a contract between the owner of an enterprise (or, more generally, a principal) and a representative worker (the agent) can specify a price at which the agent must subsequently purchase the enterprise if requested to do so by the principal. The threat of having to purchase an enterprise whose value has been diminished by his own shirking can ensure the agent puts forth the efficient level of effort, even though his effort is unverifiable. Furthermore, and perhaps more importantly, the principal can also be motivated to supply the efficient level of effort, rather than shirk and transfer the enterprise to the agent at the prenegotiated price. Consequently, the risk averse agent will never be exploited in equilibrium under the buyout scheme we propose. Nor will he bear any of the risk associated with the enterprise's uncertain profit stream.

It is the *threat* rather than the *exercise* of the buyout option that induces both parties to act efficiently in our model. As we will demonstrate, the buyout scheme we propose can completely resolve nontrivial *double* moral hazard problems, even when budget balance is required and neither delivered effort levels nor realized profits are verifiable. The right of the principal to transfer the uncertain net revenue stream to the agent if she observes (but cannot prove) he has shirked prevents him from shirking. Furthermore, knowing the agent has labored diligently, the principal does likewise and retains ownership of the firm.

This conclusion stands in contrast to results in the literature on double moral hazard. Mann and Wissink (1988), for example, ask whether money-back guarantees on a product can induce the efficient quality and maintenance decisions by a producer and consumer, respectively.[1] The authors prove the efficient decisions will generally not arise in equilibrium. We obtain a different conclusion, primarily because the agent supplies effort *before* the principal acts in our model. Thus, the principal knows how diligently the agent has toiled before undertaking any actions of her own, and this fact is exploited in the buyout option. In contrast, the producer and consumer act simultaneously in Mann and Wissink's analysis.[2]

Because our analysis documents conditions under which unverifiable information can be employed to achieve an efficient outcome in an agency setting, it is related to the recent work of Hermalin and Katz (HK) (1990). HK examine a one-sided moral hazard problem in which renegotiation of the original contract is possible. The authors focus on the setting where the agent's action is imperfectly observable to a third party (e.g., a court) in addition to being perfectly observable to the principal.[3] HK show how the threat of having the court resolve a dispute regarding the agent's effort generally enables the principal to employ her private information to discipline the agent perfectly. In our model, there is no formal renegotiation and actions are completely unobservable to third parties (e.g., courts). Furthermore, our focus is on *double* moral hazard problems, as both the worker and the owner must be motivated to act efficiently in our model.[4]

The timing, the information structure, and the potentially large wealth of the agent in our model are all special. Therefore, we do not claim the buyout option is generally applicable to all double moral hazard problems. Nevertheless, there are settings where this scheme may be relevant in mitigating nontrivial incentive problems. To illustrate, consider the interaction

between an independent retailer of a product and the product's manufacturer. The retailer can often accurately discern the quality of the goods delivered to her before she supplies the effort required to market the goods. Furthermore, the manufacturer may be a corporation with sufficient wealth to buy out the retailer. These are the key requirements for the buyout scheme we propose. Thus, our model may contribute to the literature on vertical relations between manufacturer and retailer. (See Perry (1989) for an overview.)[5]

Our analysis proceeds as follows. The model is described in section 2. In section 3, we define the first-best outcome and prove this ideal outcome can be ensured when buyout prices can be prenegotiated. The key assumptions that underlie this strong result are reviewed in section 4, where extensions of our analysis are also discussed.

2. The Model.

There are two actors in our model. Call them the "owner" and the "worker".[6] The owner is endowed with all of the bargaining power, and therefore can propose a contract to the worker on a "take-it-or-leave-it" basis.[7] The worker will agree to any contract that provides an expected utility weakly in excess of his reservation level, $\bar{U}$. $\bar{U}$ is common knowledge.

Many variables of central importance are unverifiable in the model: the agent's effort, the principal's effort, and realized net revenues. In other words, these variables cannot be observed, perfectly or imperfectly, by a third party. This forces us to consider very simple contracts. A contract can only specify: (1) a fixed payment *(P)* from the owner to the worker if the former retains title to the enterprise (i.e., does not sell the firm); and (2) a buyout price, B. The buyout price is the fixed amount the worker is obligated to pay to the owner when title to the enterprise is transferred to the worker. Such a transfer is publicly observable. The decision to transfer or

to retain ownership of the enterprise is made by the owner after the contract has been "negotiated" and after the worker and then the owner have made their contributions to the net revenues of the enterprise. Whoever holds title to the firm receives the net revenues that are generated.

We speak of each party's contribution to net revenues in terms of the level of effort supplied. The (monetary value of the) worker's effort will be denoted by $e^w \geqq 0$; the corresponding effort of the owner by $e^o \geqq 0$. Both efforts enhance expected net revenues in the enterprise. Formally, $f(R|e^w, e^o)$ is the probability density of net revenue $R \in [\underline{R}, \bar{R}]$ given effort levels e^w and e^o by the worker and owner, respectively. $F(R|e^w, e^o)$ is the corresponding distribution function. Letting subscripts denote the obvious partial derivatives, we assume $F_e(R|e^w, e^o) \leqq 0 \; \forall R, e^w, e^o$, for $e = e^w, e^o$, with strict inequality for some values of $R \in [\underline{R}, \bar{R}]$. Thus, greater effort leads to higher net revenues in the sense of first-order stochastic dominance.

The timing in the model is as follows. First, the terms of the contract (B, P) are specified. Second, the worker decides whether to accept the contract. If he accepts, the third event is the worker's choice of effort supply (e^w). Fourth, after observing the effort level supplied by the worker, the owner delivers her own effort contribution (e^o). Fifth, the owner decides whether to exercise her option to sell the enterprise to the worker. If ownership is transferred to the worker, the worker pays the owner B. If the owner retains title to the firm, she pays the worker P. Finally, net revenues are realized and accrue to the party in possession of the title to the enterprise. No recontracting between owner and worker is allowed at any time.

In this setting, the owner's decision to retain or sell the enterprise comes after effort levels have been supplied, but before net revenues are realized. Thus, the owner does not know

exactly how profitable the enterprise will ultimately be when she makes her key decision. However, her observation of the worker's effort supply provides the owner with important information about the expected value of the enterprise. This expected value is immutable once effort levels are determined. In other words, the effort decision by each party is "once-and-for-all": effort supply cannot be adjusted, *ex post*. (For example, the worker cannot initially supply a minimal level of effort and subsequently deliver more effort if he becomes the residual claimant.) This strong assumption is best viewed as characterizing the limiting case of a more general setting where effort can be supplied at two or more distinct times, but where the initial supply is most important in determining the profitability of the firm. For example, sales or repair effort, no matter how diligent, may be largely inconsequential for the profits of a firm if the firm's basic product design is seriously flawed, due to inadequate initial effort (e.g., quality provision).[8]

The agent's utility function in our model is $U(M - e^{w})$, where M denotes the monetary inflow he receives. $U(\cdot)$ is an increasing, strictly concave function of its argument. M will be equal to either: (1) P, the fixed payment delivered to the worker if the owner retains title to the enterprise; or (2) $R - B$, the realized net revenue of the enterprise less the fixed buyout price if the owner exercise her option to sell the enterprise to the worker.[9]

The owner is risk neutral. Her utility function is $T - e^{o}$, where T denotes the monetary inflow she receives. No third party can observe effort levels; but, recall, the owner privately observes the worker's effort contribution before selecting her own effort level or exercising her ownership transfer option.[10] The owner cannot make any binding commitments regarding her ultimate effort supply or ownership choice. Thus, she will make these decisions at the relevant time based solely on her own self interest at that time, given her knowledge of the preceding

history.

These timing assumptions allow a straightforward identification of equilibrium behavior.[11] At the final stage, presuming the agent has agreed to the terms of the contract, the owner will retain title to the enterprise if and only if the expected net revenues exceed the buyout price, B. This retention decision is reflected in the function $r(\cdot)$:

$$r(e^w, e^o, B, P) = \begin{cases} 1 & \text{if } \int R\, f(R|e^w, e^o)dR - P \geq B \\ 0 & \text{otherwise.} \end{cases} \tag{2.1}$$

$r(\cdot)$ is: (i) unity when the owner retains title to the firm (and pays the worker P); and (ii) zero when the title is transferred to the worker (at prespecified price B). e^w and e^o are known by the owner when she decides whether to transfer ownership.

In the penultimate stage, the compensation structure (B, P) is given, and the owner has observed the worker's effort level (e^w). The owner's maximal expected profit ($\pi^o(\cdot)$) is:

$$\pi^o(e^w, B, P) = \max_{e \geq 0} \{ r(e^w, e, B, P)[\int R\, f(R|e^w, e)dR - e - P] + [1 - r(e^w, e, B, P)][B - e] \}. \tag{2.2}$$

For later reference, denote the owner's profit-maximizing choice of effort identified in (2.2) by $\hat{e}^o(e^w, B, P)$.

Anticipating this self-interested action by the owner, the worker chooses an effort level to maximize his expected utility, $\pi^w(\cdot)$:

$$\pi^w(B, P) = \max_{e \geq 0} \{ [1 - r(e, \hat{e}^o(e, B, P), B, P)][\int U(R - e - B) f(R|e, \hat{e}^o(\cdot))dR] + r(e, \hat{e}^o(\cdot), B, P)\, U(P - e) \}. \tag{2.3}$$

At the initial stage, the agent will agree to the terms of the contract provided $\pi^{W}(B, P) \geqq \bar{U}$.

3. Findings.

The ideal outcome from the owner's perspective is the *first-best* outcome. In the first-best outcome: (1) the owner retains title to the enterprise, so $r(\cdot) = 1$; (2) the worker receives no rents, so $\pi^{W}(\cdot) = \bar{U}$; and (3) efficient effort levels, (e^{W*}, e^{O*}), are supplied by the worker and owner, respectively, where

$$(e^{W*}, e^{O*}) \in \underset{e^{W}, e^{O}}{argmax} \int R\, f(R \mid e^{W}, e^{O}) dR - e^{W} - e^{O}.^{12}$$

The first-best outcome ensures the maximum level of expected total surplus while minimizing the worker's share of this total. Expected total surplus is maximized when efficient effort levels are supplied and the risk neutral owner bears all of the risk associated with the uncertain flow of net revenues.[13]

As reported in Proposition 1, the owner can ensure this ideal outcome in the setting under consideration despite the limited set of policy instruments at her disposal.

Proposition 1. The first-best outcome can be achieved as equilibrium behavior (e^{W*}, e^{O*}) under the following compensation structure:[14]

$$B = \int R f(R \mid e^{W*}, e^{O*}) dR - e^{O*} - P; \text{ and} \tag{3.1}$$

$$U(P - e^{W*}) = \bar{U}. \tag{3.2}$$

Proof. Suppose the worker supplies effort level e^{W*}. If the owner provides effort level e, her

expected return is:

$$\text{maximum} \left\{ \int R\, f(R|e^{w*}, e)dR - P - e\,;\; B - e \right\}$$

$$\leq \text{maximum} \left\{ \int R\, f(R|e^{w*}, e^{o*})dR - P - e^{o*}\,;\; \int R\, f(R|e^{w*}, e^{o*})dR - P - e^{o*} - 0 \right\}$$

$$- \int R\, f(R|e^{w*}, e^{o*})dR - P - e^{o*} - B.$$

Thus, upon observing e^{w*}, the owner can do no better than to supply e^{o*} and subsequently retain ownership of the enterprise.

Now suppose the worker sets $e^{w} > e^{w*}$ under the identified B and P. Then the owner can secure

$$\max_{e \geq 0} \{ \int R\, f(R|e^{w}, e)dR - e - P \} > \max_{e \geq 0} \{ \int R\, f(R|e^{w*}, e)dR - e - P \} - B.$$

Hence, the owner will choose e^{o} such that $\int R\, f(R|e^{w}, e^{o})\, dR - e^{o} - P > B$, and will retain ownership. This implies $\pi^{w}(B, P) = U(P - e^{w}) < U(P - e^{w*}) = \bar{U}$.

Now suppose the worker chooses $e^{w} < e^{w*}$ under the identified B and P. Then

$$B - \max_{e \geq 0} \{ \int R\, f(R|e^{w*}, e)dR - e \} - P > \max_{e \geq 0} \{ \int R\, f(R|e^{w}, e)dR - e \} - P.$$

The owner will therefore set $\hat{e}^{o}(e^{w}, B, P) = 0$ and direct the worker to acquire the firm. Consequently, the worker faces:

$$\begin{aligned}\pi^{w}(B, P) &= \int U(R - e^{w} - B) f(R \mid e^{w}, 0) dR \\ &\leq U\left(\int R\, f(R \mid e^{w}, 0) dR - e^{w} - B\right) \\ &= U\left(\int R\, f(R \mid e^{w}, 0) dR - e^{w} - \int R\, f(R \mid e^{w*}, e^{o*}) dR + e^{o*} + P\right) \\ &< U(P - e^{w*}). \end{aligned} \tag{3.3}$$

The first inequality in (3.3) holds because the worker is averse to risk. The equality follows from (3.1). The last inequality follows from adding and subtracting e^{w*}, and noting that

$$\int R f(R \mid e^{w*}, e^{o*}) dR - e^{w*} - e^{o*} > \int R f(R \mid e^{w}, 0) dR - e^{w} \qquad \forall\, e^{w} < e^{w*}.$$

Therefore, the first-best effort levels (e^{w*}, e^{o*}) are supplied in equilibrium under the identified B and P.[15] Furthermore, with effort levels (e^{w*}, e^{o*}), $r(\cdot) = 1$ and $\pi^{w}(\cdot) = U(P - e^{w*}) = \bar{U}$ from (3.2), ensuring the first-best outcome. ■

The contract that gives rise to the first-best outcome is relatively simple. The wage payment (P) is calculated to provide the worker with exactly his reservation level of utility when he supplies the efficient level of effort, e^{w*}. (See equation (3.2).) The buyout price is set equal to the difference between: (1) the expected value of the enterprise given efficient effort levels by both the worker and owner, and (2) the sum of the worker's wage and the (value of) the owner's efficient level of effort. (See equation (3.1).) Given these specifications of B and P, it is apparent from (2.1) that the owner will retain title to the enterprise whenever the efficient effort levels have been supplied. (Notice that $\int R\, f(R \mid e^{w*}, e^{o*}) dR - P = B + e^{o*} > B$.) What may be less obvious is why efficient effort levels will be induced under the proposed contract.

As the proof of Proposition 1 makes clear, the buyout price (B) is set sufficiently high that if the worker ever puts forth less than the efficient level of effort, the owner will prefer to sell the enterprise to the worker rather than retain ownership herself and "make up for" the worker's

slack. Furthermore, since the buyout price effectively compensates the worker only for his lost wages (P) and the value of the owner's efficient level of effort (e^{O*}), the worker will prefer not to be forced to buy the risky enterprise. Consequently, he will put forth the efficient level of effort to ensure the owner finds it in her interest to supply the efficient effort level, e^{O*}, and to retain ownership of the enterprise.

Thus, the ability to specify a fixed price at which the worker must purchase the enterprise if directed to do so by the owner is of significant value to the owner. Absent this option, the owner would be unable to induce any effort from the worker in our model. With this option available, the owner can ensure the first-best outcome, even though no measure of the worker's performance is verifiable. The buyout directive provides the necessary control apparatus, even though it is not exercised in equilibrium.[16]

4. Conclusions and Extensions.

Our modest goal in this article was to demonstrate the potential value of prespecified buyout prices in resolving double moral hazard problems in environments where unverifiable information is available. In the specific setting considered, the actions of both the worker and the owner were unverifiable, as was the result of these actions (i.e., realized net revenues). Consequently, payments could only be based on the owner's decision to retain ownership of the enterprise or to sell it to the worker. We showed the ability to specify a fixed buyout price at which the worker is obligated to purchase the enterprise will enable the owner to ensure her most preferred outcome. In particular, she and the worker will always supply the efficient level of effort. Furthermore, the risk averse worker will never be expropriated; nor will he ever be forced to buy the risky enterprise in equilibrium. Thus, the double moral hazard problem is resolved completely and costlessly, even without the presence of a third party to relax the

balanced budget constraint. The unexercised buyout option provides the requisite incentives. The key feature of the buyout option is the ability to transfer ownership of the uncertain revenue stream.

The buyout option in our model plays a role similar to the role played by renegotiation in Hermalin and Katz' (1990) model. When an imperfect public signal (x) of the agent's action is available in HK's model, contracts can be conditioned on the realization of x.[17] In equilibrium, renegotiation of the original contract occurs before x is realized to avoid imposing risk on the risk-averse agent. However, the threat of enforcing the original contract allows the principal to motivate the agent and achieve the first-best outcome. In our model, the terms of the renegotiation are effectively decided in advance when P and B are specified.[18] These variables are selected to ensure the risk averse agent (the worker) bears no risk in equilibrium. P and B can also be chosen to motivate both the agent and the principal (the owner) to act efficiently in our model.

As noted, this conclusion stems from some special features of our model. The absence of limits to the worker's liability (as in Sappington (1983)) is one such feature. While an assumption of unbounded wealth for the worker is, in general, not necessary for our findings, it is important that the (off equilibrium) threat of imposing potentially large penalties on the worker be credible. If the worker does not have access to sufficient funds to buy out the owner, the scheme we propose will not be feasible. (See Romano (1991).)

The timing is another special feature of our model. It is important that the owner chooses her effort level after observing the worker's contribution.[19] If the owner and worker choose their effort levels simultaneously in the setting we analyze, the efficient effort levels (e^{w*}, e^{o*})

will not constitute Nash equilibrium behavior under any compensation structure. To see this, notice when (2.1) is imposed, e^{w*} will only be the worker's best response to e^{o*} if $\int R\, f(R|e^{w*}, e^{o*})\, dR - P \leq B$. However, e^{o*} will only be the owner's best response to e^{w*} if $\int R\, f(R|e^{w*}, e^{o*})\, dR - P \geq B + e^{o*}$. Since $e^{o*} > 0$, these two inequalities cannot hold simultaneously.

The ability of the owner to observe the worker's action is also important in our model. However, the essence of our argument carries over to some environments with imperfect monitoring. To illustrate, consider the setting where the owner is only able to observe the level of effort supplied by the worker with probability $\alpha \in (0,1)$. With the complementary probability, $1 - \alpha$, the owner sees nothing. In this environment, it can be shown that efficient effort levels (which give rise to the first-best outcome) will be supplied in equilibrium when α is strictly less than, but sufficiently close to, unity. The buyout scheme specifies a buyout price that is actuarially fair precisely when efficient effort levels have been supplied. Since the worker is strictly risk averse and the net revenues of the firm are stochastic, the worker strictly prefers not to take title to the enterprise even when efficient effort levels have been supplied. Consequently, the threat of transferring the enterprise to the worker only when he is observed to have shirked can be sufficient to deter the worker form shirking, provided shirking is observable to the owner with sufficiently high probability.[20]

Although the particular buyout scheme we have examined will not ensure the first-best outcome when the owner and worker select their effort levels simultaneously, simultaneous play does not necessarily preclude the ideal outcome. To see this, suppose the parties can commit to a randomized payment structure conditional on the owner's public report of the agent's effort

level. Further suppose the agent is paid P as defined in (3.2) if the owner reports $e^{w} \geq e^{w*}$, and $\tilde{P}$ otherwise. If $\tilde{P}$ is the realization of a random variable with mean P and sufficiently large variance, then the first-best outcome can generally be ensured if the worker is sufficiently risk averse.

This is just one of the many variations on the buyout theme that remain to be explored. The intent of this research was simply to suggest the possible importance of incentive schemes of this general form when productive inputs are observable but not verifiable.[21] Another variation of our model that warrants further study has the worker able to provide additional *ex post* effort to enhance the expected net revenues of the enterprise if he becomes the residual claimant. We conjecture the first-best outcome will continue to arise as an equilibrium provided the impact of the worker's *ex post* effort on expected profits is sufficiently small relative to the corresponding impact of the *ex ante* efforts of both owner and worker.

An environment where we suspect the first-best outcome would not be achieved so readily is one where the owner acquires private information about the expected profitability of the enterprise before deciding whether to exercise her option to sell the enterprise to the worker. In such a setting, the potential for expropriating the worker emerges, since the owner may be tempted to force the worker to purchase the enterprise when an unfavorable exogenous event occurs, even though the worker has provided the efficient level of effort.

FOOTNOTES

1. Also see the related works of Cooper and Ross (1985), Holmstrom (1982), and Khambu (1984). The latter two authors focus on the critical role a third party can play in relaxing the "balanced budget restriction" that requires total payments by and total receipts of the critical actors to be identical. (Also see Eswaran and Kotwal (1984).) As noted above, the ideal outcome is achieved in our model even with a balanced budget restriction. Romano (1991) explores a double moral hazard problem in which a manufacturer and retailer supply inputs that affect final demand. Romano's focus is on the role of resale price maintenance in a setting where prenegotiated buyout arrangements are not possible.

2. Timing is also critical in the analysis of Cremer and Riordan (1985). The authors demonstrate the gains that arise in a public goods setting when one player moves before all others. Related observations are made by Banerjee and Beggs (1989).

3. HK do present some results for the case where the third party is completely uninformed about the agent's action. (See, in particular, their Proposition 3 and the associated footnote.) We comment further on these results and provide additional comparisons between HK's work and our own in section 4.

4. There are two other papers that may be tangentially related to ours. Cramton, Gibbons and Klemperer (1987) examine conditions under which ownership of an enterprise can be assigned efficiently. Adverse selection is the focus of their model, however, while moral hazard is the focus of ours. Also, no party is able to observe the private information of any other party in their model, whereas the owner can monitor the worker's activities in

our model.

Aron (1989) develops a model in which a firm may transfer ownership of one of its divisions to the manager of that division. Such a "spinoff" in Aron's model provides more accurate information (from the securities market) about the true value of the division. The threat of acquiring the superior information facilitates better control of division managers within the firm. A transfer of ownership to the worker in our model does not provide meaningful information about either the worker's performance or the value of the firm.

5. As one other illustration, consider a partnership that provides transportation services. The first partner is responsible for vehicle maintenance, the second for training drivers. If the second partner is able to observe the first partner's diligence before supplying her training services, all moral hazard problems can be avoided with the buyout scheme we propose.

6. By treating the worker as a single entity, we abstract from the issue of how to best motivate effort from members of a team. See Alchian and Demsetz (1972), Holmstrom (1982), McAfee and McMillan (1987), and Rasmusen (1987) for insightful analyses of this interesting and difficult problem.

7. This simplifying assumption results in the owner usurping all of the expected surplus. More generally, an *ex ante* bargaining stage could be appended to our model. This stage could serve to divide between the owner and the worker the (maximal) expected surplus that is generated by the incentive scheme described below.

8. Additional thoughts on the extent to which this assumption can be relaxed are presented in section 4.

9. We do not allow the worker to contract with a third party. In particular, we do not consider a sale of the newly-acquired enterprise by the worker to another party.

10. In section 4, we provide some thoughts on the case of imperfect monitoring, where the owner cannot always observe the worker's performance. Also see Demski and Sappington (1991) and Hermalin and Katz (1990).

11. We concentrate on the subgame played between the owner and worker, once the contract terms are agreed upon. Capitalizing on the information structure and sequential nature of the subgame, we then focus on subgame perfect equilibria.

12. In particular, efficient effort levels are characterized by

$$\int R\, f_{e^W}(R|e^{W*}, e^{O*})\, dR = \int R\, f_{e^O}(R|e^{W*}, e^{O*})\, dR = 1.$$

e^{W*} and e^{O*} are assumed to be strictly positive and finite. For simplicity, both are also assumed to be unique. Standard Inada and concavity conditions will guarantee these properties hold.

13. Notice the implicit assumption that production is profitable in the first-best outcome.

14. As the proof of Proposition 1 makes clear, the identified equilibrium behavior is unique if, when indifferent among choices, the owner always acts to maximize total expected surplus $(\pi^W(\cdot) + \pi^0(\cdot))$.

15. Notice this conclusion holds even if the agent is risk neutral. With a risk neutral agent, the first inequality in (3.3) would be replaced by an equality. But the second inequality would remain strict, ensuring the critical result. Strict risk aversion only becomes important in a setting where the principal's observation of the agent's effort is not always perfect. (See the discussion in section 4.)

16. Our buyout scheme is somewhat reminiscent of the scapegoat scheme Rasmusen (1987) proposed to resolve moral hazard problems in teams. The scapegoat scheme imposes a severe penalty on a randomly-selected team member (the scapegoat) if the observed performance of the team falls short of the specified goal. Under some conditions, fear of being selected as the scapegoat can induce all strictly risk averse members of the team to supply their efficient, unobservable effort levels. The fear of being forced to bear an unfavorable lottery also motivates the worker in our model. In our model, though, the production function is not deterministic, and the ultimate performance of the enterprise is not verifiable. Furthermore, the owner is given the right to designate whether she or the worker will serve as the "scapegoat" (i.e., own the enterprise) in our model. Of primary concern in our analysis is ensuring the owner will not abuse this right in equilibrium.

17. HK point out in a footnote that in settings where the enterprise can be sold (as it can in our model), there need be no observable realization of x. All that is needed in this case is for the sale of the enterprise to be observable.

18. We wish to thank Benjamin Hermalin for suggesting this interpretation of our buyout scheme.

19. Of course, the owner could make her effort supply decision after the retention decision, or simultaneously with the rentention decision.

20. Although this extension of our basic model introduces imperfect monitoring of the worker by the owner, it does so in a special way. After monitoring, the owner either knows exactly what the worker has done, or she knows nothing about the worker's activities. An alternative monitoring technology would be one that reported the agent's actual level of effort with error. See Demski and Sappington (1991) and Hermalin and Katz (1990).

21. Of course, the buyout arrangement presented here will implement the first-best outcome in a standard principal - agent model where the principal supplies no productive input. Our conclusions also extend to the case where the agent's preferences are separable (as they are assumed to be in Hermalin and Katz (1990)), provided standard concavity conditions are satisfied.

REFERENCES

Alchian, A. and Demsetz, H. "Production, Information Costs, and Economic Organization." *American Economic Review*, Vol. 62 (1972), pp. 777-795.

Aron, D. "Corporate Spinoffs in an Agency Framework." Northwestern University MEDS discussion paper, revised May 1989.

Banerjee, A. and Beggs, A. "Efficiency in Hierarchies: Implementing the First-Best Solution by Sequential Actions." *Rand Journal of Economics*, Vol. 20 (1989), pp. 637-645.

Cooper, R. and Ross, T. "Product Warranties and Double Moral Hazard." *Rand Journal of Economics*, Vol. 16 (1985), pp. 103-113.

Cramton, P., Gibbons, R. and Klemperer, P. "Dissolving a Partnership Efficiently." *Econometrica*, Vol. 55 (1987), pp. 615-632.

Cremer, J. and Riordan, M. "A Sequential Solution to the Public Goods Problem." *Econometrica*, Vol. 53 (1985), pp. 77-84.

Demski, J. and Sappington, D. "Incentive Compatible Use of Unverifiable Information." Yale University mimeo, February 1991.

Eswaran, M. and Kotwal, A. "The Moral Hazard of Budget-Breaking." *Rand Journal of Economics*, Vol. 15 (1984), pp. 578-581.

Hermalin, B. and Katz, M. "Moral Hazard and Verifiability." University of California - Berkeley mimeo, May 1990.

Holmstrom, B. "Moral Hazard in Teams." *Bell Journal of Economics*, Vol. 13 (1982), pp. 324-340.

Khambu, J. "Optimal Product Quality Under Asymmetric Information and Moral Hazard." *Bell Journal of Economics*, Vol. 13 (1982), pp. 483-492.

Mann, D. and Wissink, J. "Money-Back Contracts with Double Moral Hazard." *Rand Journal of Economics*, Vol. 19 (1988), pp. 285-292.

McAfee, R. P. and McMillan, J. "Optimal Contracts for Teams." University of Western Ontario mimeo, August 1987.

Perry, M. "Vertical Integration: Determinants and Effects." In *The Handbook of Industrial Organization*, R. Schmalensee and R. Willig (eds.), Amsterdam: North-Holland, 1989.

Rasmusen, E. "Moral Hazard in Risk-Averse Teams." *Rand Journal of Economics*, Vol. 18 (1987), pp. 428-435.

Romano. R. "Double Moral Hazard and Resale Price Maintenance." University of Florida mimeo, January 1991.

Sappington, D. "Limited Liability Contracts Between Principal and Agent." *Journal of Economic Theory*, Vol. 29 (1983), pp. 1-21.