Working Paper 9118

RISK AVERSION, PERFORMANCE PAY, AND THE PRINCIPAL-AGENT PROBLEM

by Joseph G. Haubrich

Joseph G. Haubrich is an economic advisor at the Federal Reserve Bank of Cleveland. The author thanks Linda Fuhr for typing assistance.

December 1991

ABSTRACT

This paper calculates numerical solutions to the principal-agent problem and compares the results to the stylized facts of CEO compensation. The numerical predictions come from parameterizing the models of Grossman and Hart and of Holmstrom and Milgrom. While the correct incentives for a CEO can greatly enhance a firm's performance, providing such incentives need not be expensive. For many parameter values, CEO compensation need only increase by about $10 for every $1,000 of additional shareholder value; for some values, the amount is 0.003 cents. The paper thus answers two challenges posed by Jensen: that principal-agent theory does not yield quantitative predictions, and that CEO compensation is insufficiently sensitive to firm performance.

I. Introduction

The principal-agent paradigm lies at the center of corporate finance. Its central problem of motivating a subordinate provides sharp insight into issues surrounding debt, equity, dividends, executive pay, and regulators' activities during the thrift debacle. Jensen and Murphy (1990) challenge this paradigm, finding that the compensation of top executives increases a mere $3.25 per $1,000 gain in shareholder wealth. This pay-performance sensitivity of 0.003 is a far cry from the 1.0 predicted by the risk-neutral version of principal-agent theory. Their estimate challenges broader versions of the theory to predict anything quantitative at all. In moving beyond risk neutrality, "...theory says nothing about the magnitude of the pay/performance relation" (Baker, Jensen, and Murphy [1988], p. 611).

Yet, as this paper shows, principal-agent theory can yield exact quantitative predictions. Grossman and Hart (1983) provide such a solution for a two-state, finite-action model. Once a few parameters are chosen, quantitative results follow. Thus, it becomes possible to find what principal-agent theory predicts for the cases Jensen and Murphy consider.

In the parameterized models, small amounts of risk aversion can result in quite low values of pay sensitivity. The results in this paper thus respond to both of the challenges issued by Jensen and Murphy: providing quantitative predictions, and predicting low pay-performance sensitivity. The results also confirm their intuition that incentives matter greatly for executive performance and consequently for shareholder value.

One motive for pursuing the quantitative approach is the success it has found in related fields. Asset pricing has long benefited from a fruitful interplay between empirical work and quantitative models. So has public finance, where authors from Mirrlees (1971) to Sheshinski (1989) have used quantitative solutions to the "hidden type" (as opposed to the "hidden action" type considered here) principal-agent problem to calculate the optimal income tax.

The following section provides more detail on pay sensitivity. Section III reviews the model and the needed analytical results, all of which follow from Grossman and Hart. Section IV reports the numerical solutions, and section V compares the results with similar calculations based on a model from Holmstrom and Milgrom (1987). Section VI concludes.

II. Performance Pay

The argument about pay sensitivity has both a descriptive and a prescriptive component. In the descriptive part of their analysis, Jensen and Murphy (1990) carefully gather data and assess the pay-performance sensitivity for chief executive officers (CEOs) in the Forbes Executive Compensation Survey, encompassing salary, bonus, stock options, insider stock holdings, and dismissal probability. By the narrowest measure, salary plus bonus, which excludes savings plans, thrift plans, other benefits, and stock holdings and options, the pay-performance sensitivity is 0.0000135, or 1.35 cents per $1,000 of shareholder value. By the broadest measure, which looks at changes in CEO wealth and includes Forbes' "total compensation," stock options and insider stock holdings, present value of changes in salary and bonus, and the effects of dismissal, total pay-performance sensitivity is $3.25 per $1,000, or 0.00325.

The prescriptive part is emphasized more in Jensen's Harvard Business Review article, "Eclipse of the Public Corporation" (1989). He argues here that publicly traded corporations do not resolve the conflict between shareholders and managers as well as do leveraged buyout (LBO) associations. In LBO associations, executive salaries are close to 20 times more sensitive to performance than in public corporations. As Jensen states (p. 69), "It's not hard to understand why an executive who receives $200 for every $1,000 increase in shareholder wealth will unlock more value than an executive who receives $3.25."

In addition to suggesting that theory cannot explain the quantitative fact, Jensen and Murphy also present qualitative evidence against a standard principal-agent interpretation for executive pay. They point out the superior performance of LBO associations, the decline in equity holding among officers over time, and the sparse use of additional information in compensation schemes. The principal-agent model can potentially explain these results, because parameter values can differ across firms and across time. This paper generally avoids those issues, though section V briefly considers why additional information may be ignored.

III. Model and Analytical Results

Though the principal-agent model is a natural one for executive compensation, with shareholders as the principal and the CEO as the agent, most versions have intractable solutions. Quantitative predictions do not readily emerge from the implicit equations that define the sharing rules, especially in models with a continuum of states. Grossman and Hart (1983, sec. 4), however, provide a tractable solution to the two-state, finite-action case. Choosing a few parameters and functional forms then results in a quantitative prediction.

To appreciate both the scope and the limitations of the numerical calculations, it helps to know the basic model structure and the theorems that justify the later work.

Assumptions and Notation

The principal, or firm, has two gross profit levels, $q_1<q_2$. The agent can choose from a finite set A of actions, $\{a_1,a_2,...a_n\}$, which influences the probability of good and bad profits. These might be viewed as effort levels, or different projects the CEO can approve. Formally, $\Pi_i(a)$ denotes the probability of state (profit level) i given action a.

The agent's utility depends on actions and on income, expressed as U(a,I). Solving the model requires some restrictions on U, expressed in

Assumption A1: $U(a,I)= G(a)+K(a)V(I)$, where

i) V is real, continuous, strictly increasing, and concave on an open interval (I_ℓ, ∞);

ii) $\lim_{I\Rightarrow I_\ell} V(I)=-\infty$;

iii) G and K are real, continuous on A and $K>0$;

iv) For all a,b contained in A and all I,J contained in (I_ℓ,∞),

$G(a)+K(a)V(I)\geq G(b)+K(b)V(I)$ implies

$G(a)+K(a)V(J)\geq G(b)+K(b)V(J)$.

This formulation makes preferences over income risk independent of action (the converse is also true), and makes the agent risk averse.

The numerical calculations of section IV further specialize the utility function, setting either G=0 or K=1. One useful function of this form is the constant absolute risk aversion (CARA) type with

$V(I) = -e^{-\delta I}$, $K(a) = e^{\delta a}$, so $U(a,I) = -e^{-\delta(I-a)}$,

in which effort appears as negative income. The disutility of effort greatly influences management compensation; treating effort as negative income makes the resulting contract easier to interpret.

The agent also has a $\overline{U}$, derived from an alternative employment or leisure-time activity he can turn to if he does not work for the principal. I exclude actions so distasteful that the agent would never choose them, and to keep the problem interesting I assume that some actions remain. This <u>participation constraint</u> is embodied in

<u>Assumption A2</u>: For every action a contained in A, there is an I in (I_ℓ, ∞) such that $G(a)+K(a)V(I) \geq \overline{U}$.

The final assumption puts some structure on the probability of profits.

<u>Assumption A3</u>: For all a contained in A and i=1...n, $\Pi_i(a) > 0$.

This rules out the Mirrlees (1974) plan of imposing increasingly high penalties with increasingly low probabilities as the agent takes actions approaching the right one: Since $\Pi_i(a) > 0$ and A is finite, such a strategy puts too much risk on the agent.

One major advance of Grossman and Hart is to focus on the cost of getting the agent to choose a particular action. In the "first best" case, in which the principal observes the action, the cost is simply the agent's reservation price for a, denoted

$$CFB(a) = h[(\overline{U} - G(a))/K(a)], \text{ where } h = V^{-1}.$$

Of course, the whole point of the principal-agent problem is that the principal cannot observe the action. She can only make payment dependent on the realized output state, the gross profit level. This incentive scheme, a set of payments contingent on the state $\{I_1, I_2\}$, gives the agent utility levels

$v_1 = V(I_1)$

$v_2 = V(I_2)$.

Although the principal cannot observe the action, she can design an incentive scheme that induces the agent to choose that action. This defines the second-best cost of an action a*, C(a*).

(1) Choose v_1, v_2 to minimize $\sum_{i=1}^{2} \Pi_i (a*) h \quad v_i)$

subject to

$$G(a*) + K(a*) \, [\sum_{i=1}^{2} \Pi_i (a*) \, v_i] \geq G(a) + K(a) \, [\sum_{i=1}^{2} \Pi_i (a) v_i] \forall a \epsilon A$$

$$G(a*) + K(a*) \, [\sum_{i=1}^{2} \Pi_i (a*) \, v_i] \geq \overline{U}$$

$$v_i \epsilon \; Range \; (V) \; \forall i.$$

The first constraint is the incentive compatibility constraint, which states that the agent takes action a* only if that action gives a higher payoff than any other action. The second and third constraints are the participation constraints, which state that the agent must get a certain minimum utility, and that there exists an income level that produces the desired utility.

Several incentive schemes (I or v sets) may induce the agent to choose action a* (that is, implement a*). Define C(a*) as the greatest lower bound (infinum or inf) of $\Sigma \Pi_i \; h(v_i)$ on the constraint set of equation (1). For an empty constraint set, set C(a*) to infinity. In that case, the principal cannot induce action a*.

A little terminology about the principal completes the basic notation. Because the risk-neutral principal gets the gross profits, her expected benefit from an agent's action is

$$B(a) = \sum \Pi_i (a) q_i.$$

The optimal action (second-best) maximizes the expected net benefit to the principal,

$$B(a)-C(a).$$

General Results

The simple solution used below depends on some general theorems due to Grossman and Hart. At least a passing familiarity with these ideas is necessary in order to understand the range and restrictions on the results.

From the general n-state case, I take two main results. Proposition 1 states that assuming A1, A2, and A3, there exists a second-best optimal action and a second-best optimal incentive scheme. Proposition 6 states that for finite A, the agent is indifferent between the action he takes and some less costly actions under the optimal incentive scheme.

Three results for the two-state case make possible an explicit calculation of the solution. First, a definition: Action a is efficient if the probability of a good outcome can be increased only by incurring a greater cost. Proposition 10 states that assuming A1, A2, A3, and V are strictly concave, with two states every second-best action is efficient.

The next propositions also restrict possible solutions. Proposition 11 states that under the conditions of Proposition 10, the agent obtains his reservation utility. Proposition 12 states that assuming A1, A2, and two states, adding new actions a', such that CFB(a') $\geq$ CFB(a) for all a in A, cannot make the principal worse off. That is, adding distasteful actions won't encourage the agent to shirk but might give the principal more information.

Solution Techniques for the Two-State, Finite-Action Case

Grossman and Hart have a simple plan for solving the principal-agent problem. First, compute the cost C(a) for each action a. Then optimize the net benefit, B(a)-C(a), over all actions a. Several special results make this unusually easy in the two-state, finite-action case.

Proposition 6, that the agent is indifferent between the chosen action and some less costly action, combines with Proposition 10, that the agent chooses only efficient actions, to drastically simplify the cost and probability structure. Without loss of generality, we can assume that $CFB(a_1)<CFB(a_2)<...<CFB(a_n)$ and correspondingly $\Pi_2(a_1) < \Pi_2(a_2) < ...< \Pi_2(a_n)$. Because there is no conflict in getting the agent to take the least costly (minimum effort) action a_1, $C(a_1)=CFB(a_1)$. This ties down $C(a_1)$.

For the other actions, I use Proposition 6 and Proposition 11 to compute $C(a_k)$. For each action a_j for j<k, find the I_1, I_2 pair that makes the agent indifferent between a_k and a_j and that also sets that common expected utility to $\overline{U}$. Finding v_1 and v_2, the utility levels that satisfy those two conditions, involves solving two simultaneous linear equations.

$$(2) \qquad G(a_k) + K(a_k) [\Pi_1 (a_k) v_1 + \Pi_2 (a_k) v_2] = \overline{U}$$

$$G(a_j) + K(a_j) [\Pi_1 (a_j) v_1 + \Pi_2 (a_j) v_2] = \bar{U}.$$

Solving for v_1 and v_2 yields

$$(3) \qquad v_1 = \frac{\Pi_2 (a_j) \left[\frac{\bar{U} - G(a_k)}{K(a_k)}\right] - \Pi_2 (a_j) \left[\frac{\bar{U} - G(a_j)}{K(a_j)}\right]}{\Pi_1 (a_k) - \Pi_1 (a_j)}$$

$$v_2 = \frac{\Pi_1 (a_j) \left[\frac{\bar{U} - G(a_k)}{K(a_k)}\right] - \Pi_1 (a_k) \left[\frac{\bar{U} - G(a_j)}{K(a_j)}\right]}{\Pi_2 (a_k) - \Pi_2 (a_j)}.$$

The incentive scheme, or actual payments to the agent, comes from inverting the utility function, setting $I_1 = h(v_1)$ and $I_2 = h(v_2)$.

This gives (k-1) different (v_1, v_2) pairs and hence (k-1) different (I_1, I_2) pairs, one for each $j<k$. Proposition 6 says that one of these pairs must be the minimum cost-incentive scheme for a_k; that will tell us $C(a_k)$. Grossman and Hart show that incentive compatibility implies it is the pair with the largest v_2. With any other pair, the agent would prefer action a_j to action a_k.

The (v_1, v_2) pair chosen for each action a_k must then be checked against the range of V. For example, exponential or power utility functions are always negative, but the above procedure sometimes demands strictly positive utility. In that case, $C(a_k)$ is infinite: The principal cannot induce that behavior from the agent. For the feasible v_i's, cost is simply the expected value of payments to the agent, or

$C(a_k) = \Pi_1(a_k)v_1 + \Pi_2(a_k)v_2$.

Once the second-best cost $C(a_k)$ is computed for each k, the problem becomes straightforward. The principal chooses among a finite number of actions to maximize B(a)-C(a).

IV. Numerical Solutions

Section III outlines a way to calculate exact quantitative solutions to simple principal-agent models. I now choose the parameters and perform the calculations, applying the model to the executive compensation problems of Jensen and Murphy.

The model has many free parameters. These include the payoff, risk aversion, disutility of effort, reservation utility, and outcome state probabilities, as well as the effects of action on probabilities, number of actions, and functional forms. Even those parameters previously estimated, such as risk aversion, do not have standard, accepted values. Others, such as the CEO's effect on share value, are conceivably measurable, but serious practical problems prevent measurement. Still others, such as the number of actions, have no real empirical counterpart.

To overcome this, I present a variety of solutions for different parameter values. When possible, I use the estimates of Jensen and Murphy, such as the standard deviation of shareholder value and the average compensation of CEOs. For parameters they do not estimate, such as disutility of effort, I take values from their illustrative examples. For the remaining parameters, such as risk aversion, I use a range of values. Thus, the model predicts outcomes for the cases Jensen and Murphy discuss using parameter values that are close to what most people would consider sensible.

Basing the parameter values on the Jensen and Murphy examples represents only a small subset of possible predictions. Some later examples explore the broader range of possibilities by using more extreme values.

Base Case

The first (base) case has two states and three actions. The standard deviation of shareholder wealth in the Jensen and Murphy samples is $200 million (p. 244). For a two-point distribution with x > y, the standard deviation is $\sqrt{P(1-P)}$ (x-y). For p=1/2, this further simplifies to ($\frac{1}{2}$) (x-y). Putting all monetary rewards in units of $1 million ($\$10^6$), I choose gross profit levels of 300 and 700.

Jensen and Murphy do not specify a reservation wage; I use the median CEO compensation for their sample, $490,000. For later robustness checks, I use $250,000, which is Business Week's estimate of Sanford Grossman's academic salary (Byrne [1988]).

I use constant absolute risk aversion (CARA), or exponential utility. This method has several advantages: It provides a simple interpretation of effort as negative income, it is robust to differences in the wealth of the agents, and it facilitates comparison with the recent consumption literature (Caballero [1990]), which also uses CARA. Unfortunately, it makes comparisons with the asset pricing literature, which uses constant relative risk aversion (CRRA), more difficult. Given wealth levels, though, it is easy to calculate relative risk aversion. As an additional check, I provide a few calculations using CRRA.

More parameters come from recasting an example from Jensen and Murphy

(p. 228). The CEO considers a project he privately values at $100,000, but which costs the firm $10 million. This translates into the principal-agent framework as follows: Let the set of acts be A={0.001,0.1,0.2}. To the agent, taking action a_1 is like paying $1,000, taking action a_2 is like paying $100,000, and taking action a_3 is like paying $200,000. From the shareholder's viewpoint, each action increases the probability of the good state, where shareholder value is 700, by 0.025, so

$\Pi_2(a_1) = 0.475$

$\Pi_2(a_2) = 0.5$

$\Pi_2(a_3) = 0.525$.

Taking action a_2 instead of a_1 thus results in (0.025)x(700-300), or a $10 million gain in expected shareholder value at a personal cost of $99,000. I choose nine values for risk aversion, setting $\gamma = 0.125$ to 1.125 by increments of 0.125. Because γ measures absolute risk aversion, it is not strictly comparable to the more common measures of relative risk aversion, such as the 29 of Kandel and Stambaugh (1991) or the 2-3 of Friend and Blume (1977). Using a wealth estimate from Jensen and Murphy (CEO-controlled company stock holdings) of $8.8 million gives a relative risk aversion of 1.1 to 9.9.

Figure 1 presents the results for the base case, plotting the profit share against risk aversion. Profit share is the increase in CEO compensation divided by the increase in shareholder value between the good and the bad state, or Wage(state2)-Wage(state1)/(700-300). Table 1 reports the actual numbers. The profit share fraction, when positive, varies from around 0.01 to 0.03, implying that CEOs get an extra $10 to $30 for a $1,000 increase in shareholder value. Table 1 reports an overabundance of negative wages. Realistically, negative compensation sometimes results, because stock holdings represent such a large share of total compensation. Requiring only positive wages, while equally unrealistic, would dramatically alter the principal-agent problem (Sappington [1983]).

For $\gamma > 0.625$, the profit share is zero, implying the agent takes the easiest possible action. Income shows no response to shareholder value: The profit share fraction is 0. Everyone knew this held as γ approached infinity, but it is a definite surprise that a number less than 1 counts as "close to infinity." For some plausible parameter values, the 0.00325 value of Jensen and Murphy looks too sensitive. Perhaps the low profit share Jensen and Murphy find in the data reflects an average including a few zeros.

The positive profit shares in table 1 underscore a related message of Jensen, however. The response of executive pay to firm performance, though slight, significantly increases the firm's value. A positive profit share makes sense only if it induces the agent to work more, to choose a_2 or a_3 instead of a_1. In my example, this is a gain to shareholders of $10 million or $20 million. Because the principal must compensate the agent for the risk involved, if the agent still takes the lowest action he gains nothing and the principal loses. Furthermore, because the principal (shareholders) maximizes net benefits, a less sensitive scheme (though cheaper) would induce the agent to pick a lower action, at a substantial cost.

Do these results make sense? Is there any reason behind the particular values shown in table 1? One advantage of an explicit model is that we can explore such questions and deepen our intuition about the problem. In general, the compensation contract represents a trade-off between insurance and incentives. A risk-neutral agent would bear all of the risk and accept a profit share of 1, but a risk-averse agent would naturally desire to shift some of the output risk to the risk-neutral principal. Using the model, we can quantify both sides of the trade-off and understand what motivates the principal and the agent.

The contracts in table 1 clearly share risk between the principal and the agent. At a risk aversion of 0.125, wages in the good and bad state differ by $4 million (four units in the table). The agent would pay $365,000 to avoid this risk. To keep the agent at his reservation utility, the principal must compensate him with a risk premium for accepting this uncertainty. If the agent bore the full profit uncertainty (±$200 million), he would demand a risk premium of $194 million. As Sappington ([1991], p. 49) puts it, "To conserve on the risk premium she must award the agent for bearing risk, the principal will choose to bear some risk herself."

Providing insurance to the agent creates its own problems, however. As noted by Sappington (1991, pp. 49-50), "When he is effectively insured against bad outcomes under the optimal contract, the agent will exert less effort to avoid these bad outcomes." A closer look at the base-case contracts of table 1 shows that they do provide strong incentives to the agent despite relatively low profit shares.

One reason behind the low profit share is the difference between expected gains and realized gains. For a risk aversion of 0.125, the wage difference of $4 million translates into a rather modest profit share 0.01, or $10 per $1,000. But the CEO does not directly determine shareholder value; he merely changes the odds. Choosing a_2 over a_1 increases <u>expected</u> shareholder value by $10 million; choosing a_3 adds another $10 million. From this perspective, his compensation more closely matches his contribution.

From another perspective, the agent's pay also matches his contribution.

By taking an action that improves expected shareholder value, he increases his own expected value of pay by 0.025 x \$4,000,000, or \$100,000, his disutility of effort from taking that action.

Robustness Checks

The base case directly confronts the theoretical challenge posed by Jensen and Murphy, but as a single example, its results might be special or unrepresentative. The next few tables and figures report on variations of the base case. Changes include risk aversion and the agent's effect on profit, number of acts, and reservation utility.

The first set varies both risk aversion and the agent's effect on the probability of the gross profit level. The relation between the agent's effort and the probability of success is given by another new parameter, APROB. As APROB gets bigger, good actions increase the expected probability of the good state by an increasing amount. In the three-act case,

$\Pi_2(a_1) = 0.5\text{-APROB}$

$\Pi_2(a_2) = 0.5$

$\Pi_2(a_3) = 0.5\text{+APROB}.$

In the next two tables, APROB varies from 0.00625 to 0.05625 by increments of 0.00625. It thus provides a range around the base-case value of 0.025.

Table 2 and figure 2 report the results. For high APROB and low risk aversion, table 2 shows CEOs getting between \$4 and \$5 per \$1,000, moving a step closer to the \$3.25 finding of Jensen and Murphy.

Figure 3 presents the results with a reservation wage of 0.25 (\$250,000). The share results hardly change, as befits CARA. Wage levels are correspondingly lower, but are not reported.

Figure 4 and table 3 report the results for a calculation using 11 acts instead of the three used above, with similar results.

To some people, the relative risk aversion (γ times wealth) in these examples may seem too high. The next two figures address that problem. Figure 5 looks at absolute risk aversion between 0.0025 and 0.0225. Figure 6 shows profit share for very low levels of risk aversion. Even a billionaire would have relative risk aversion below 0.00000001 with these preferences. Even so, the pay-performance ratio stays close to 0.01, or \$10 for every \$1,000. Low performance pay does not require high risk aversion.

In addition, figures 5 and 6 together emphasize the nonlinearity and nonmonotonic relationship between risk aversion and profit share. They show both a local minimum and a more extreme nonmonotonicity.

Extreme Cases

The principal-agent model predicts a wider range of behavior than the results so far suggest. The profit share can approach one quite closely. It can also approach zero and still retain some pay-performance sensitivity.

Figure 7 shows the results of changing A to {0.01,10,20}, APROB to 0.025, and the risk-aversion range to 1×10^{-12} to 9×10^{-12}. The figure shows nonlinearity and a profit share around 0.999, or \$999 per \$1,000.

Figure 8 and Table 4 report the results for the CRRA case. This calculation holds risk aversion constant at 9 and varies APROB from 0.05 to 0.45, in increments of 0.05. This means $\Pi_2(a_3)$ varies from 0.55 to 0.95. The act set is A={1,10,15}, and the disutility of effort function is $(1/20)a^2$. The reservation wage is 0.25. These parameters result in a low profit share fraction, which ranges from 0.000000033 to 0.0000003. This implies that executive compensation should increase between 0.003 and 0.03 cents per $1,000 of shareholder wealth. In some cases, then, it takes little to motivate a CEO.

V. Comparison with Holmstrom and Milgrom Results

Holmstrom and Milgrom (1987), by making stronger assumptions, simplify the calculations even more. They take a strategy diametrically opposed to that behind the two-state, finite-action case. They start by complicating the model. Essentially, the agent has so much freedom that only simple linear rules provide the correct incentives.

The agent controls the drift rate of a stochastic process z over a time period of one unit. The process z evolves according to the equation $dz=\mu dt + dB$. The agent has a CARA utility, with quadratic disutility of effort. The cost to the agent of controlling the drift is $c(\mu)=(k/2)\sigma^2$. Holmstrom and Milgrom show that the optimal compensation rule takes the form $\alpha z + \beta$, with the optimal pay-performance ratio of

(4) $$\alpha* = \frac{1}{1+k\ \gamma\sigma^2}.$$

Under this incentive scheme, the agent chooses a constant value of μ. In that sense, the problem reduces to a static setting.

I translate the Jensen and Murphy examples into this framework as follows. First, let the unit time period be one year. The controlled process is shareholder wealth, whose standard deviation of 200 means a variance of 40,000. Assume a risk aversion γ of one. If a project worth $10,000,000 (10 units) to the shareholders decreases the utility of the agent by $100,000 (0.1

units), this implies a k value of 0.018. Substituting these values into equation (4), I obtain a pay-performance sensitivity of 0.00138, or $1.38 for every $1,000.

More generally, the formula expressing profit share as a function of risk aversion, keeping the disutility at 0.1 for 10, is $1/(1+800\ \gamma\ e^{-\gamma/10})$.

Figure 9 plots profit share against risk aversion for 100 values of risk aversion running from 0.025 to 2.5. Profit share runs from $47 per $1,000 of shareholder value to 6 cents per $1,000.

The Holmstrom-Milgrom approach has less flexibility than that of Grossman and Hart. It can neither use CRRA nor measure effort as negative income, much less use the more general forms allowed by A1. Controlling a drift (but not variance) term is less general than changing outcome probabilities. For these reasons, the Grossman-Hart approach seems the more desirable one.

The Holmstrom-Milgrom approach, however, has an additional advantage beyond computational simplicity. It can explain why the principal ignores additional information about the agent's action. For example, shareholders may have information about revenues and expenses, in addition to profits. When the agent has some discretion over how to account for revenues and expenses, this information should not influence his compensation. Note further that even in the simplest case embodied in equation (4), a great deal of information is ignored. Agent compensation depends only on a time aggregate of total profits at year end, not on performance at each date.

At this level of analysis, there is no need to choose between the two approaches. The two sets of results confirm each other. Both predict low levels of pay-performance sensitivity, even for low levels of risk aversion.

VI. Conclusion

Owners and managers must decide how to share profits. That decision lies behind executive compensation, shareholding, the firm's debt-equity mix, and even takeover policy. Yet, the quantitative theoretical predictions about that sharing are rare: a share of 1 for a risk-neutral agent, and a share of 0 for an infinitely risked agent. Where do real-world cases fit in?

This paper has used two methodologies, one due to Grossman and Hart and one due to Holmstrom and Milgrom, to compute quantitative solutions. The results should not be considered a test of principal-agent theory. The paper does not formally confront a specified hypothesis with data. It should rather be seen as an application of that theory.

A test is, in principle, possible with this methodology. Assuming a distribution across parameters will produce a distribution across profit shares and wages, and this distribution could be compared with estimated values. Empirically identifying the model requires restricting the joint distribution of APROB, reservation utility, and risk aversion, about which I have little intuition. It seems better to leave explicit tests for future work.

Still, I would like to claim that the results do provide information about the correspondence between principal-agent theory and reality. Two lessons emerge from the exercise: (1) Low profit shares can occur with low risk aversion, and (2) even low profit shares provide incentives and substantially increase the value of the firm.

REFERENCES

Baker, George P., Michael C. Jensen, and Kevin J. Murphy. "Compensation and Incentives: Practice vs. Theory," Journal of Finance, vol. 43, July 1988, pp. 593-613.

Byrne, John A. "It's So Hard to Keep Good Professors Nowadays," Business Week, November 28, 1988, p. 92.

Caballero, Ricardo J. "Consumption Puzzles and Precautionary Savings," Journal of Monetary Economics, vol. 25, no. 1, January 1990, pp. 113-136.

Friend, Irwin, and Marshall E. Blume, "The Demand for Risky Assets," American Economic Review, vol. 65, 1977, pp. 900-922.

Grossman, Sanford J., and Oliver D. Hart. "An Analysis of the Principal-Agent Problem," Econometrica, vol. 51, no. 1, January 1983, pp. 7-45.

Holmstrom, Bengt, and Paul Milgrom. "Aggregation and Linearity in the Provision of Intertemporal Incentives," Econometrica, vol. 55, no. 2, March 1987, pp. 303-328.

Jensen, Michael C. "Eclipse of the Public Corporation," Harvard Business Review, vol 67, September/October 1989, pp. 61-74.

_______, and Kevin J. Murphy. "Performance Pay and Top-Management Incentives," Journal of Political Economy, vol. 98, no. 2, April 1990, pp. 225-264.

Kandel, Shmuel, and Robert F. Stambaugh. "Asset Returns and Intertemporal Preferences," Journal of Monetary Economics, vol. 27, no. 1, February 1991, pp. 39-71.

Mirrlees, James. "An Exploration in the Theory of Optimum Income Taxation," Review of Economic Studies, vol. 38, 1971, pp. 171-208.

_________. "Notes on Welfare Economics, Information, and Uncertainty," in Essays on Economic Behavior under Uncertainty, ed. G. M. Balch, D. McFadden, and Shik-Yen Wu, North Holland, Amsterdam, 1974.

Sappington, David. "Limited Liability Contracts between Principal and Agent," Journal of Economic Theory, vol. 29, 1983, pp. 1-21.

_________. "Incentives in Principal-Agent Relationships," Journal of Economic Perspectives, vol. 5, no. 2, Spring 1991, pp. 45-66.

Sheshinski, Eytan. "Note on the Shape of the Optimum Income Tax Schedule," Journal of Public Economics, 1989, pp. 201-215.

Table 1

Wages and Profit Shares - Base Case

Risk Aversion	Wage: Bad State	Wage: Good State	Profit Share Fraction
0.125	-1.187	2.878	0.010
0.250	-1.026	3.341	0.0109
0.375	-0.897	4.242	0.012
0.500	-0.792	10.88	0.029
0.625	0.491	0.491	0.00
0.750	0.491	0.491	0.00
0.875	0.491	0.491	0.00
1.000	0.491	0.491	0.00
1.125	0.491	0.491	0.00

Source: Author's calculations.

Table 2

Wages and Profit Shares

Panel A. Wages

Risk Aversion	**0.125**		**0.250**		**0.375**		**0.500**		**0.625**	
Probability Gain	Bad State	Good State	Bad State	Good State	Bad State	Good State	Bad State	Good State	Bad State	Good State
0.00625	0.491	0.491	0.491	0.491	0.491	0.491	0.491	0.491	0.491	0.491
0.0125	-2.64	6.09	-2.17	21.6	0.491	0.491	0.491	0.491	0.491	0.491
0.188	-1.70	3.81	-1.45	4.94	-1.25	14.5	0.491	0.491	0.491	0.491
0.0250	-1.19	2.88	-1.03	3.34	0.897	4.24	-0.792	10.9	0.491	0.491
0.0313	-0.862	2.36	-0.751	2.62	-0.659	3.01	-0.582	3.76	-0.517	8.70
0.0375	-0.638	2.04	-0.556	2.20	-0.488	2.42	-0.429	2.76	-0.378	3.40
0.0438	-0.473	1.82	-0.411	1.93	-0.358	2.07	-0.312	2.27	-0.271	2.57
0.0500	-0.348	1.65	-0.299	1.73	-0.257	1.83	-0.219	1.96	-0.186	2.14
0.0563	-0.249	1.53	-0.210	1.59	-0.175	1.66	-0.144	1.75	-0.116	1.87

Risk Aversion	**0.750**		**0.875**		**1.000**		**1.13**	
Probability Gain	Bad State	Good State	Bad State	Good State	Bad State	Good State	Bad State	Good State
0.00625	0.491	0.491	0.491	0.491	0.491	0.491	0.491	0.491
0.0125	0.491	0.491	0.491	0.491	0.491	0.491	0.491	0.491
0.188	0.491	0.491	0.491	0.491	0.491	0.491	0.491	0.491
0.0250	0.491	0.491	0.491	0.491	0.491	0.491	0.491	0.491
0.0313	0.491	0.491	0.491	0.491	0.491	0.491	0.491	0.491
0.0375	-0.333	7.23	0.491	0.491	0.491	0.491	0.491	0.491
0.0438	-0.234	3.13	-0.202	6.18	0.491	0.491	0.491	0.491
0.0500	-0.155	2.14	-0.128	2.90	-0.103	5.38	0.491	0.491
0.0563	-0.0907	2.03	-0.0676	2.28	-0.0465	2.72	-0.0271	4.76

Source: Author's calculations.

Table 2

Wages and Profit Shares

Panel B. Profit Shares

Risk Aversion	0.125	0.250	0.375	0.500	0.625	0.750	0.875	1.00	1.13
Probability Gain									
0.00625	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.0125	0.0219	0.0595	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.188	0.0138	0.0160	0.0393	0.00	0.00	0.00	0.00	0.00	0.00
0.0250	0.0102	0.0109	0.0128	0.0292	0.00	0.00	0.00	0.00	0.00
0.0313	0.00807	0.00842	0.00917	0.0109	0.230	0.00	0.00	0.00	0.00
0.0375	0.00669	0.00689	0.00728	0.00798	0.00946	0.0189	0.00	0.00	0.00
0.0438	0.00572	0.00585	0.00607	0.00645	0.00709	0.00841	0.0159	0.00	0.00
0.0500	0.00500	0.00508	0.00522	0.00545	0.00581	0.00641	0.00758	0.0137	0.00
0.0563	0.00444	0.00450	0.00459	0.00475	0.00497	0.00531	0.00586	0.00691	0.0120

Source: Author's calculations.

Table 3

Wages and Profit Shares

11 Act Case

Panel A: Wages

Risk Aversion	0.125		0.250		0.375		0.500		0.625	
Probability Gain	Bad State	Good State	Bad State	Good State	Bad State	Good State	Bad State	Good State	Bad State	Good State
0.00625	0.251	0.251	0.251	0.251	0.251	0.251	0.251	0.251	0.251	0.251
0.0125	-2.89	5.88	0.251	0.251	0.251	0.251	0.251	0.251	0.251	0.251
0.188	-1.95	3.59	-1.69	4.72	-1.49	15.0	0.251	0.251	0.251	0.251
0.0250	-1.43	2.65	-1.27	3.11	-1.14	4.01	-1.03	10.7	0.251	0.251
0.0313	-1.11	2.13	-0.994	2.39	-0.901	2.78	-0.823	3.52	-0.757	8.46
0.0375	-0.882	1.81	-0.799	1.97	-0.729	2.19	-0.669	2.52	-0.618	3.16
0.0438	-0.717	1.58	-0.654	1.69	-0.599	1.83	-0.552	2.03	-0.511	2.33
0.0500	-0.592	1.42	-0.541	1.50	-0.498	1.60	-0.459	1.72	-0.426	1.90
0.0563	-0.492	1.29	-0.452	1.35	-0.416	1.42	-0.384	1.51	-0.356	1.63

Risk Aversion	0.750		0.875		1.000		1.13	
Probability Gain	Bad State	Good State	Bad State	Good State	Bad State	Good State	Bad State	Good State
0.00625	0.251	0.251	0.251	0.251	0.251	0.251	0.251	0.251
0.0125	0.251	0.251	0.251	0.251	0.251	0.251	0.251	0.251
0.188	0.251	0.251	0.251	0.251	0.251	0.251	0.251	0.251
0.0250	0.251	0.251	0.251	0.251	0.251	0.251	0.251	0.251
0.0313	0.251	0.251	0.251	0.251	0.251	0.251	0.251	0.251
0.0375	-0.573	6.99	2.51	0.251	0.251	0.251	0.251	0.251
0.0438	-0.474	2.89	-0.442	5.94	0.251	0.251	0.251	0.251
0.0500	-0.395	2.17	-0.368	2.66	-0.343	5.14	0.251	0.251
0.0563	-0.331	1.79	-0.308	2.04	-0.287	2.48	-0.267	4.52

Source: Author's calculations.

Table 3

Wages and Profit Shares

11 Act Case

Panel B: Profit Shares

Risk Aversion	**0.125**	**0.250**	**0.375**	**0.500**	**0.625**	**0.750**	**0.875**	**1.00**	**1.13**
Probability Gain									
0.00625	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.0125	0.0219	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.188	0.0138	0.0160	0.0412	0.00	0.00	0.00	0.00	0.00	0.00
0.0250	0.0102	0.0110	0.0129	0.0293	0.00	0.00	0.00	0.00	0.00
0.0313	0.00810	0.00845	0.00919	0.0109	0.0230	0.00	0.00	0.00	0.00
0.0375	0.00672	0.00691	0.00729	0.00798	0.00946	0.0189	0.00	0.00	0.00
0.0438	0.00575	0.00586	0.00608	0.00645	0.00709	0.00841	0.0159	0.00	0.00
0.0500	0.00502	0.00509	0.00523	0.0546	0.00581	0.00641	0.00758	0.0137	0.00
0.0563	0.00446	0.00451	0.00460	0.00475	0.00497	0.00531	0.00586	0.00691	0.0120

Source: Author's calculations.

Table 4

Wages and Profit Shares - CRRA Case

APROB	Wage: Bad State	Wage: Good State	Profit Share Fraction
0.05	0.24994522	0.25006446	2.98E-007
0.10	0.24997498	0.25003459	1.49E-007
0.15	0.24998490	0.25002465	9.94E-008
0.20	0.24998987	0.25001968	7.45E-008
0.25	0.24999285	0.25001669	5.96E-008
0.30	0.24999483	0.25001471	4.97E-008
0.35	0.24999625	0.25001329	4.26E-008
0.40	0.24999732	0.25001222	3.73E-008
0.45	0.24999815	0.25001139	3.31E-008

Source: Author's calculations.

Figure 1

Profit Share vs. Risk Aversion—Base Case

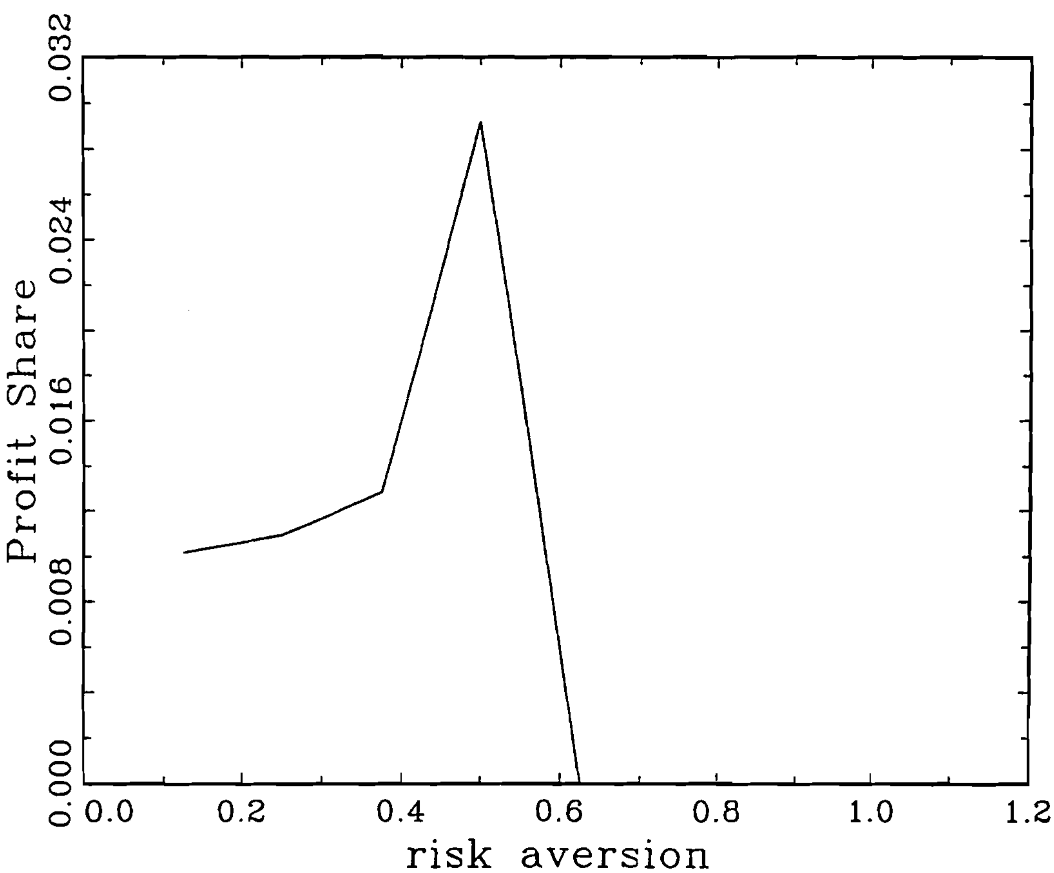

Source: Author's calculations.

Figure 2
Profit Share vs. Agent Effect and Risk Aversion

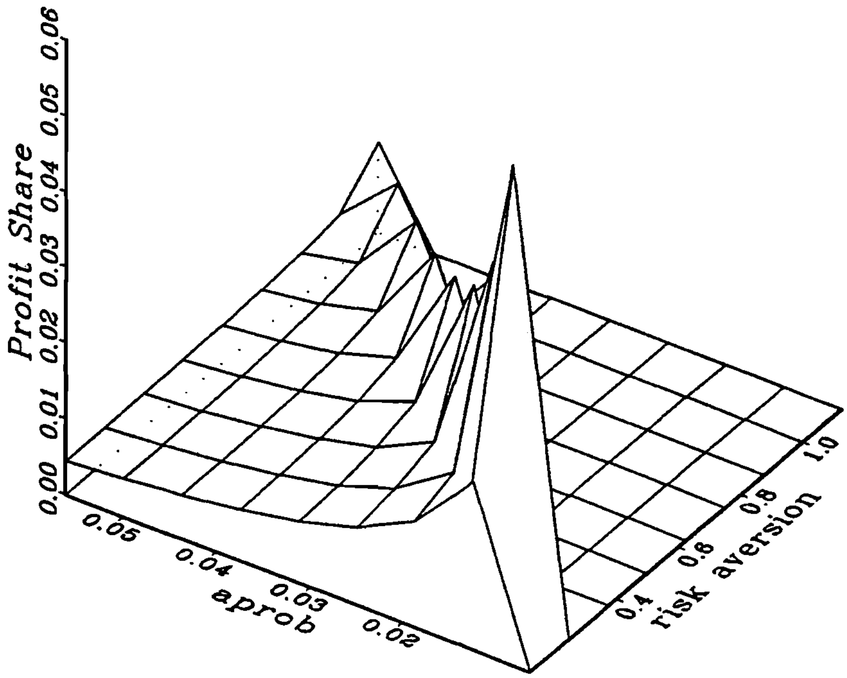

Source: Author's calculations.

Figure 3

Profit Share vs. Agent Effect and Risk Aversion—$250,000 Reservation Wage

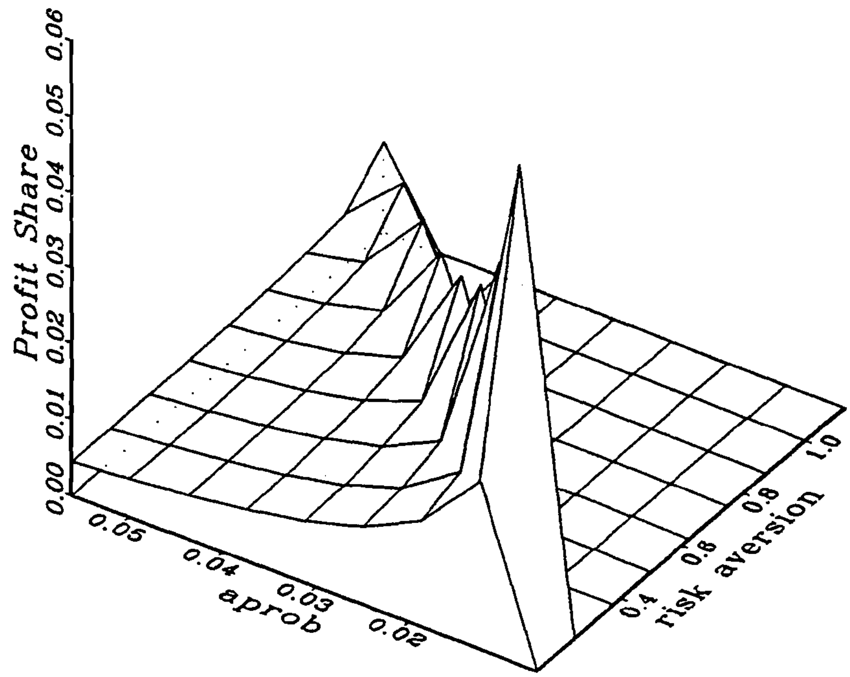

Source: Author's calculations.

Figure 4
Profit Share vs. Agent Effect and Risk Aversion—11-Act Case

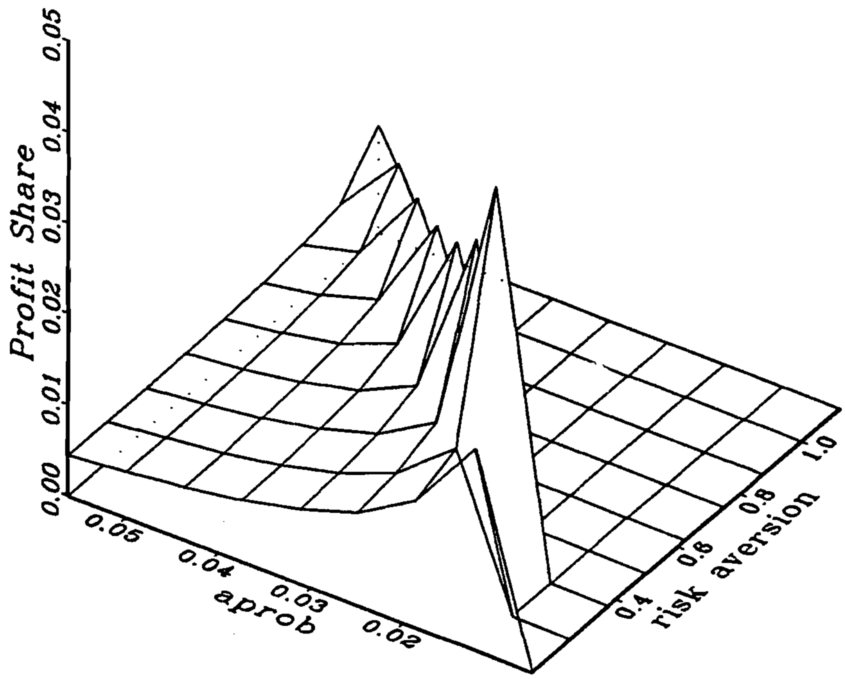

Source: Author's calculations.

Figure 5

Profit Share vs. Risk Aversion—Low Risk Aversion

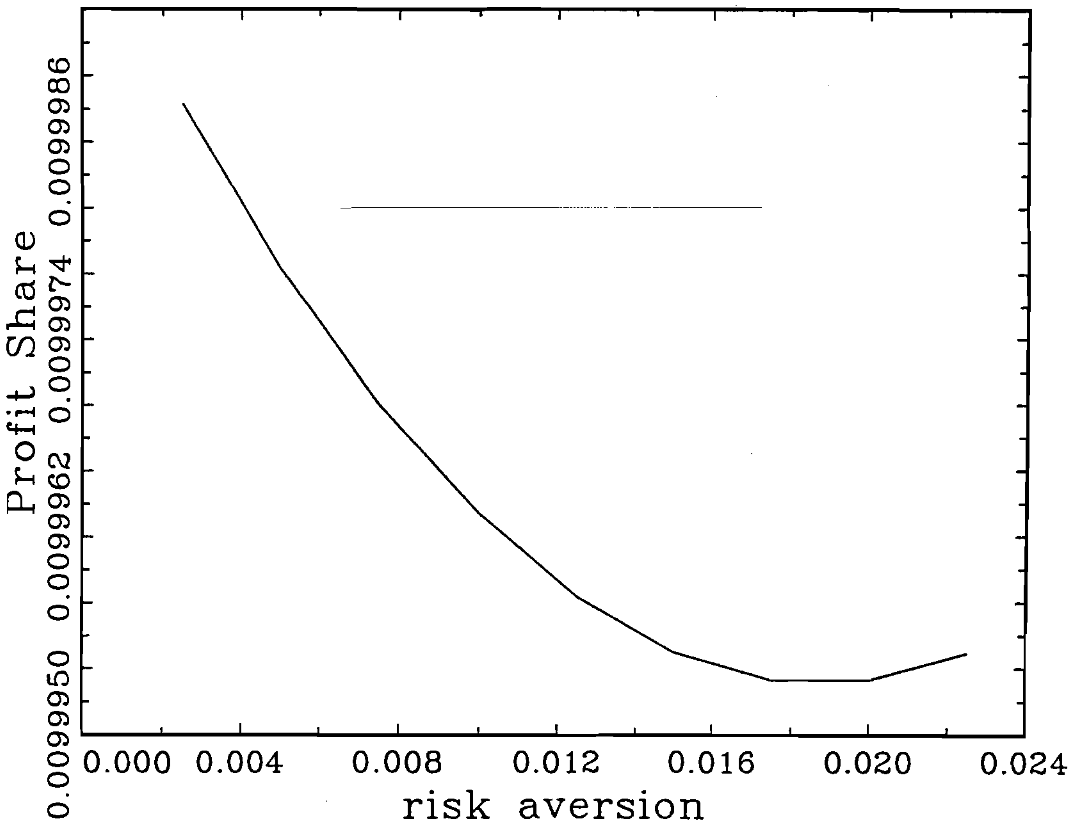

Source: Author's calculations.

Figure 6

Profit Share vs. Risk Aversion—Extremely Low Risk Aversion

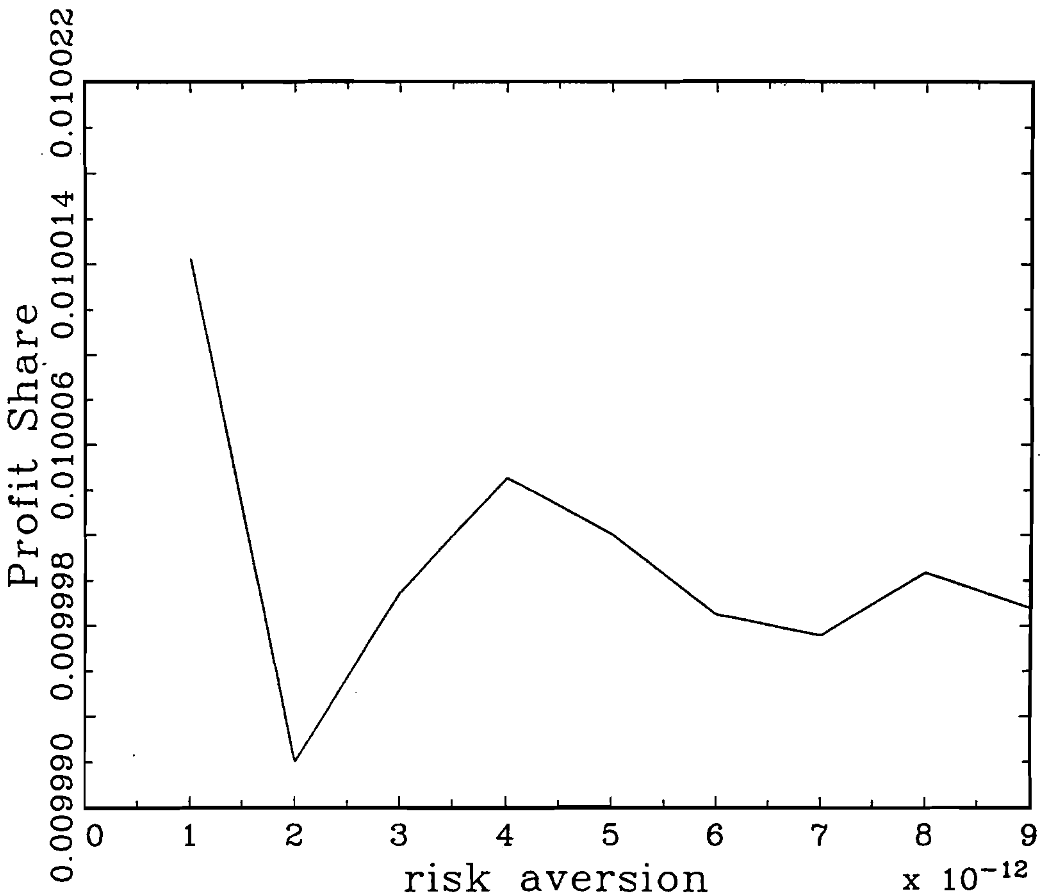

Source: Author's calculations.

Figure 7
Profit Share vs. Risk Aversion—High Profit Share

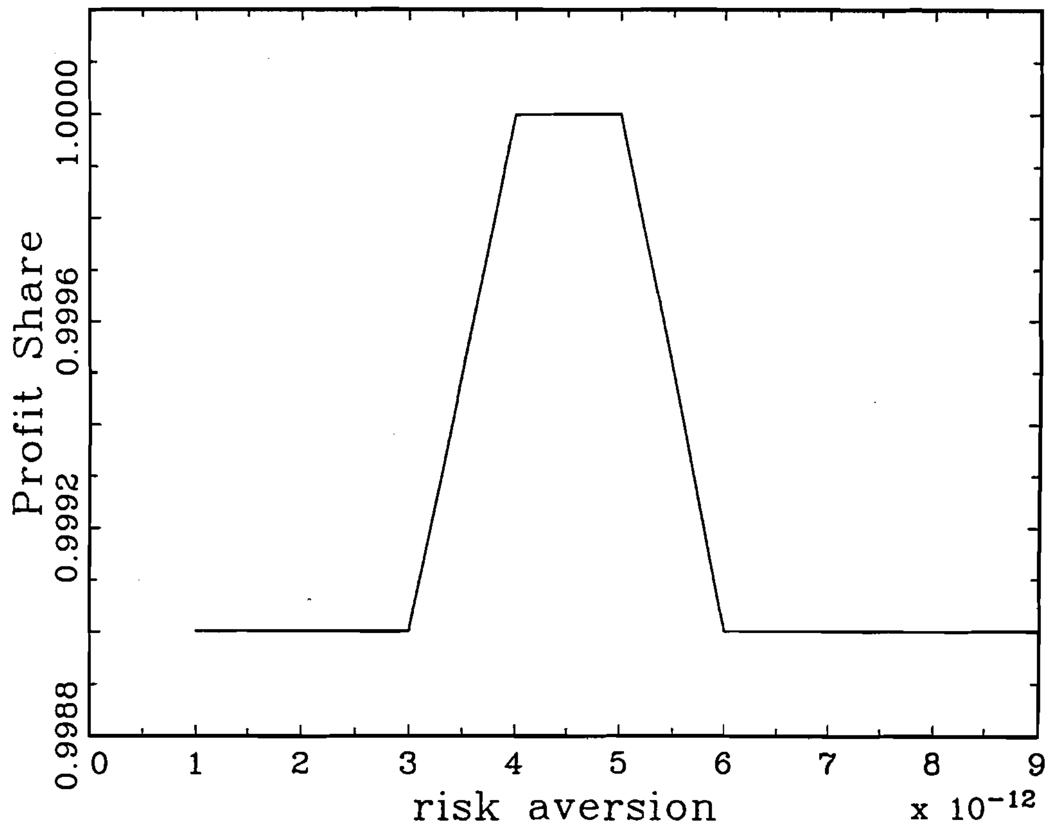

Source: Author's calculations.

Figure 8
Profit Share vs. Agent Effect—CRRA Case

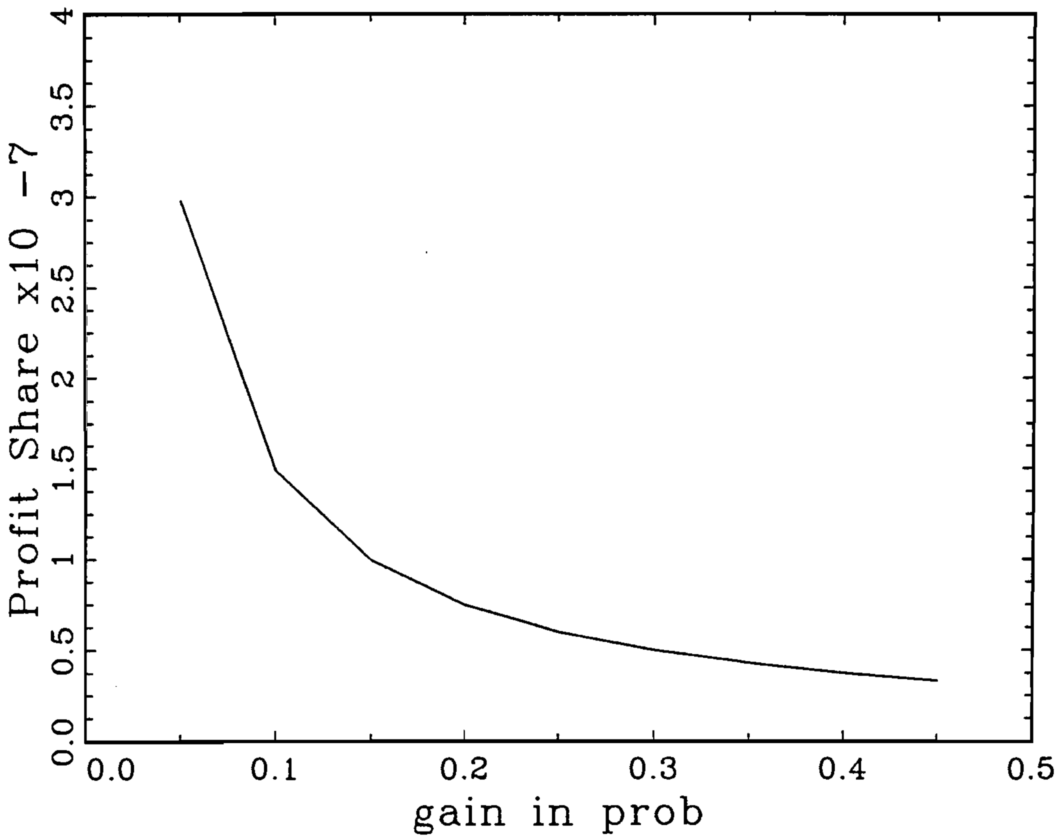

Source: Author's calculations.

Figure 9
Profit Share vs. Risk Aversion—Holmstrom-Milgrom Model

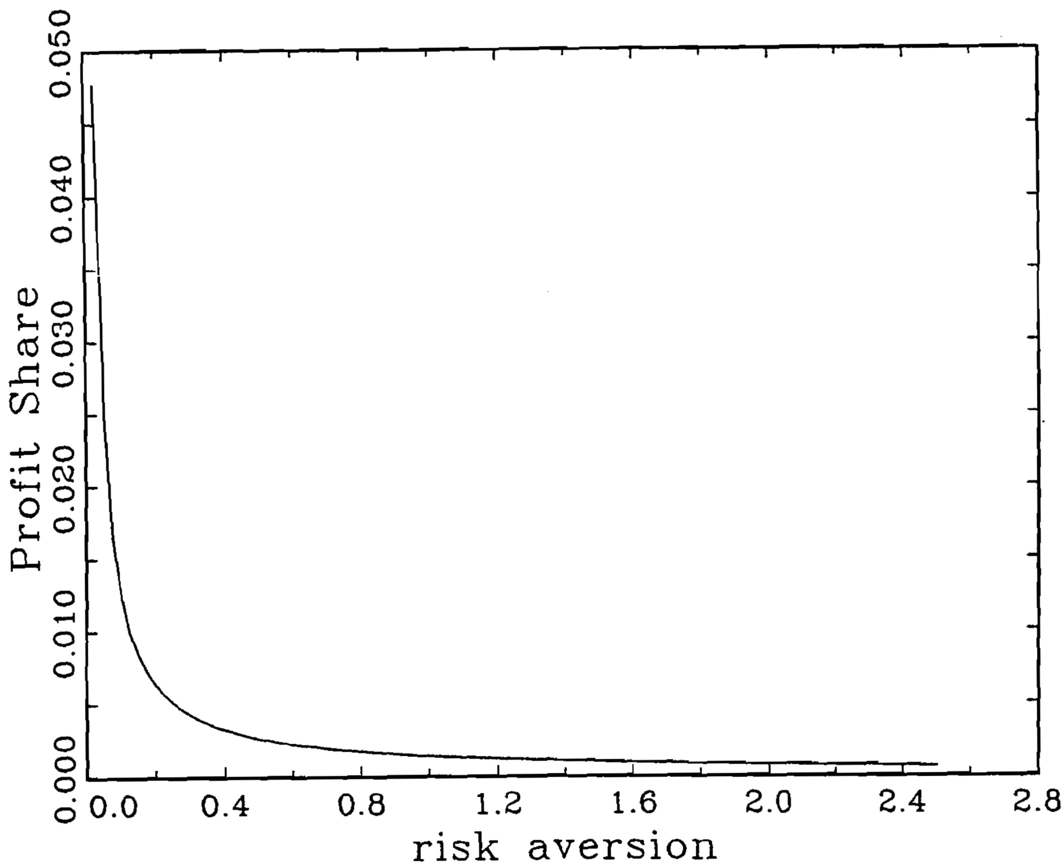

Source: Author's calculations.