

UC Berkeley

CUDARE Working Papers

Title

Managerial reputation and the "endgame"

Permalink

<https://escholarship.org/uc/item/7np2x55r>

Authors

Berck, Peter
Lipow, Jonathan

Publication Date

1999-05-01

CUDARE


DEPARTMENT OF AGRICULTURAL AND RESOURCE ECONOMICS/
DIVISION OF AGRICULTURE AND NATURAL RESOURCES
UNIVERSITY OF CALIFORNIA AT BERKELEY.

Working Paper No. 884

MANAGERIAL REPUTATION AND THE "ENDGAME"

by

Peter Berck and Jonathan Lipow



California Agricultural Experiment Station
Giannini Foundation of Agricultural Economics
May, 1999



MANAGERIAL REPUTATION AND THE “ENDGAME”

by

Peter Berck
Department of Agricultural and Resource Economics
University of California at Berkeley

Jonathan Lipow
Department of Agricultural Economics
Hebrew University of Jerusalem

ABSTRACT

KEYWORDS

Corresponding author:

Peter Berck
Department of Agricultural
and Resource Economics
207 Giannini Hall
University of California
Berkeley, CA 94720

July 17, 1997

MANAGERIAL REPUTATION AND THE “ENDGAME”

1. INTRODUCTION

A typical basketball game is characterized by an ever-widening divergence in tactics as the game approaches its conclusion (in an attempt to avoid turnovers and deny their opponents the time required to close the lead). The team that is leading tends to play cautiously and slow down the game's pace, even under the pressure of the “shot clock.” The trailing team invariably adopts the antithesis of the leading team's approach. The endgame tactics of trailing teams are dominated by fast breaks and desperate attempts for three-point baskets.

It is our belief that the incentives often facing managers concerned with enhancing or protecting their reputations are analogous to those faced in basketball's endgame. A manager is “ahead” when his performance to date has enhanced his reputation. Such a manager can be expected to avoid taking risks that may endanger his reputation even if these risks are well justified from the owners' perspective. A manager is “behind,” however, when performance to date has eroded his reputation. Such a manager must restore his reputation or face dismissal. Such managers are likely to take excessive risks in hope of salvaging their reputations. After all, the money being gambled with belongs to someone else and, if nothing is done, unemployment is unavoidable. Hence, there is little to lose.

This paper presents a model in which managers of varying abilities choose strategies for their companies while companies only decide whether or not to retain a manager. A strategy determines the mean and variance of the company's performance. Better managers have more strategies to choose from than do less capable managers. One possible Bayes-perfect equilibrium is that the less capable managers deliberately choose high-variance strategies, exactly like the trailing basketball team.

Endgame reputational incentives have implications for the firm's capital structure as well. Distortions of the firm's choices regarding risks and return may raise the cost of debt capital providing an additional constraint on the firm's attempts to attain an optimal debt/equity ratio. In situations where bankruptcy costs are high, it is even conceivable that the incentive for poor managers to take excessive risks may lead to complete extinction of certain classes of firms which, absent the moral hazard of endgame incentives, would have played a productive role in society.

In this paper, we will present a model that induces equilibrium manager behavior similar to that seen in basketball's endgame. We believe that the conditions required to induce such equilibria are often observed in the real world. Managers who are candidates for endgame-type behavior are likely to work for firms where direct observation of managers is costly leading owners to infer both manager abilities and decisions based on observation of easily identified benchmarks, such as earnings, sales, or free cash flow. Such conditions are common in firms where (1) ownership is sufficiently dispersed that the costs of monitoring manager decision making is prohibitive for any single or small group of shareholders or (2) the firm is sufficiently small that it has not attracted any objective analytical following among financial firms (obviously, analytical reports published by the firm's underwriters, who have an interest in maintaining good relations with firm managements, are not likely to be objective). Publicly-owned firms that share these characteristics are well represented on all the major stock exchanges.

A vivid illustration of endgame behavior among managers (in this case, portfolio managers) is furnished by the tale of a hapless Chilean copper trader. Working the graveyard shift, the trader incorrectly entered a trade and lost a few million dollars for Chile's national copper firm. Desiring to cover up his embarrassing error, the trader proceeded to engage in a series of futures speculations using the firm's money. The trader's original aim was to make good the initial loss before he

was reprimanded. After a series of additional losses, the trader's objective became to make good the losses before he was dismissed. As losses mounted even further, the objective became to make good the losses before being arrested. As this process developed, the level of risk taken by the trader grew ever larger. His losses were finally noticed, and the trader was arrested but only after he had managed to lose \$200 million (living proof that individuals can, indeed, have a noticeable impact on national accounts).

Given other recent financial debacles (Barings, Sumitomo), it appears that the monitoring of manager decision making is particularly costly in financial trading and that endgame behavior is rampant in financial firms. Empirical evidence that portfolio managers alter the risk/return characteristics of their investments in order to affect their reputations (as measured by the flow of money into funds that they manage) is provided in Chevalier and Ellison (1995) and Falkenstein (1996).

There are seven sections in this paper. In section 2, we briefly review the research conducted so far on the importance of managerial reputation in influencing firm decision making. In section 3, we define the basic parameters of a labor market and delineate formal mathematical conditions for Bayesian-perfect equilibria in that market. In section 4, we present a class of graphically and algebraically tractable examples. In section 5, we identify and characterize the classes of strategy sets that are consistent with the equilibrium. We demonstrate that equilibria exist which may lead some or all managers to choose inefficient mean/variance combinations (in which mean return has been sacrificed to *increase* variance). In section 6, we consider the welfare implications of changes in the distribution and quality level of managers. We conclude the paper by summarizing and analyzing our results and by discussing the implications of endgame-type behavior for mechanism design.

2. LITERATURE REVIEW

Since the subject was first introduced in Fama (1980), there has been a growing appreciation of the influence of reputational effects in dictating the behavior of agents. Holmstrom (1982) first discussed the possibility that the managers of firms, fearing that they would be revealed as inferior, could choose to forgo investment projects that owners would have found desirable. Holmstrom suggested that reputational effects could thus be used as a justification for the widely accepted belief that large firms were managed too conservatively. Holmstrom and Ricart i Costa (1986) extended Holmstrom's presentation, further emphasizing the tendency of managers to avoid desirable investments that could expose them to undesirable reputational effects. In addition, they suggested contractual mechanisms that alleviated the misalignment of owner and manager risk preferences. Contractual and informational issues related to this framework were further developed in Ricart i Costa (1989). Gibbons and Murphy (1992) test the form of managerial incentives with a sample of chief executive officers.

An extreme form of managerial conservatism is the herding behavior described in Sharfstein and Stein (1990). In their formulation, portfolio managers converged on identical strategies in order to assure that they could do no worse than average. While doing better than average carried rewards, these were outweighed by the costs of underperformance. Hence, all managers mimicked each other in order to assure that they would not be average. While herding is clearly a conservative strategy, it does not imply that clients are always exposed to less than optimal levels of risk. One of the portfolio managers quoted by Scharfstein and Stein recounts that he was well aware of the stock market's excessive risk in September of 1907 but would not lower his exposure since no one else was doing so.

Huddart (1996) presented a reputational portfolio management model that induced all types of portfolio managers to take excessive risks. In Huddart's model, one investment security's risk/return profile stochastically dominated the other. There were two portfolio managers. One was better informed than the other and, if he demonstrated this during the first period, he would be rewarded during the second period. The informed portfolio manager would receive a private signal regarding the inferior security that could make it more attractive. If the information was sufficiently favorable, he would overweight the inferior security. Although all managers chose their portfolios simultaneously, the uninformed manager's need to maximize his chance of appearing to be the informed manager would lead him to overweight the inferior security as well even though his information did not justify such a decision. Meanwhile, the informed manager's desire to maximize his chance of appearing to be informed would lead him to overweight the inferior security to an extent greater than that justified by the superior information that he possessed.

Our model is closest to Zwiebel (1995). He also addressed managerial conservatism, presenting a model in which managers had two alternative investment projects to choose from. One project return profile stochastically dominated the other, but the inferior project's outcome more clearly signaled the manager's true level of ability. In Zwiebel's formulation, average managers preferred the inferior investment that clearly signaled that they were average, exceptionally capable managers chose the superior investment since they were confident that their abilities would be recognized anyway, and poor managers chose the superior investment since they counted on the noisier signal to, perhaps, mask their true level of ability.

Our model differs from Zwiebel in permitting both managers and firms a wider range of options while restricting the number of types of agents to two. We do not consider optimal incentive contracts. With the expended choices, the firm and managers are players in a game whose equilibrium concept is Bayes-perfect. In this

expanded framework, there are many more types of equilibria than in Zwiebel, including the intentional choice by less capable managers of strategies that lead to excessive variance in the firm's performance. Thus, the model presented below can generate both reckless and conservative behavior in equilibrium.

3. A GENERAL TWO-PERIOD MODEL

In this model, an owner hires a manager, observes the manager's performance in the first of two periods, and then decides whether to retain or replace the manager prior to the second period. Managers may either be good or bad. The owner knows the population frequency of both types of manager, but has no way of telling whether a given manager is good or bad. Therefore, she hires a manager at random. The manager that has been hired takes an action, σ , which, along with a random process, determines an outcome, Y , for the first period. The owner, who benefits from the outcome Y , infers from the realized value of Y the likelihood that the manager is good. Based on this inference, the owner decides whether to retain the manager for a second period or fire him and hire a new manager at random for the second period. The set of realizations of Y that lead the owner to replace the manager is known as C —the critical region for the owner's test of the manager's abilities. The set of realizations of Y that lead the owner to retain the manager is known as \bar{C} . A graphic depiction of the flow of events in this model is presented in Figure 1.

3.1. The Nature and Behavior of Managers

A manager hired by the owner will always receive a one-period contract with a fixed and nonnegotiable payment. Managers will always prefer being employed by the owner to their next best alternative employment. As a result, the manager hired by the owner prior to the first period will make choices that maximize his chance of keeping his job for a second period.

Managers are not identical but are divided into two types (good, G, and bad, B) that differ in their ability to generate Y; Y is a random variable with likelihood function $l(y, \mu, \sigma)$, where μ is the mean of Y and σ is Y's standard deviation. Mean μ is a function of the manager's type as well as of the value of σ . Managers will be able to unobservably choose the value of σ . For now, μ and σ will be sufficient statistics to characterize Y (later in this paper, we introduce the assumption that Y is distributed normally).

A manager of type i 's ability to generate Y is limited to combinations bounded by a continuous and differentiable mean/variance frontier $\mu_i(\sigma)$. The frontier is comprised of three segments: an "efficient" segment (where $d\mu/d\sigma > 0$), an "inefficient" segment (where $d\mu/d\sigma < 0$), and a transition point (where $d\mu/d\sigma = 0$ and μ is at its maximum). Both μ and σ must always be greater than or equal to zero.

Inefficient mean/variance combinations are generally ignored in the economics/finance literature, since it is assumed that no one will ever desire to choose such a combination. As we show below, there may, indeed, be situations in which inefficient combinations may be chosen by managers.

The mean/variance frontier of a good manager strictly dominates that of a bad manager. A graphic representation of mean/variance frontiers attainable by good and bad managers is shown in Figure 2.

For a given owner's critical region, the manager will choose the mean/variance combination that maximizes his chance of retention. Let $l(y|\mu_G(\sigma), \sigma)$ and $l(y|\mu_B(\sigma), \sigma)$ be the likelihood functions for Y given the manager's type and choice of σ . Since C is the critical region, the optimal choices of σ for both types of manager, σ_G^* and σ_B^* , will be given by the values of σ that minimize the probability of a realization of Y inside of C (i.e., that minimizes the probability of being fired):

$$(3.1a) \quad \sigma_G^* = \operatorname{argmin}_{y \in C} \int l(y|\mu_G(\sigma), \sigma) dy$$

$$(3.1b) \quad \sigma_B^* = \operatorname{argmin}_{y \in C} \int l(y|\mu_B(\sigma), \sigma) dy.$$

The managers' choices of variance, σ_G^* and σ_B^* , are the best replies to the owner's choice of C .

The manager hired by the owner for the second period will not be concerned with keeping his job for a third period, since there are only two periods in the model. For simplicity, we will assume that the manager hired by the owner for the second period, unable to alter his own prospects, makes choices that maximize the owner's expected utility. This assumption may be relaxed. As long as good managers always make second-period choices that give the owner greater expected utility than the choices made by bad managers, the basic results of this paper are not affected.

3.2. *The Behavior of the Owner*

For a given pair of managers' choices for variance, σ_G^* and σ_B^* , the owner will choose a set, C , that maximizes her expected utility over both periods. Since the owner chooses a manager at random for the first period, first-period expected utility must equal

$$(3.2) \quad P_G \cdot \int u(y) l(y|\mu_G(\sigma_G^*), \sigma_G^*) dy + P_B \cdot \int u(y) l(y|\mu_B(\sigma_B^*), \sigma_B^*) dy,$$

where $u(y)$ is the owner's utility given outcome y , P_G is the probability that a manager chosen at random is good, and P_B (which equals $1 - P_G$) is the probability that a manager is bad.

Let the utility that the owner expects to enjoy in the second period if the manager is type (G) be $EU(Y|G)$ while $EU(Y|B)$ is the utility that she expects to

enjoy if the manager is type (B). Let P_{1G} equal the owner's Bayesian posterior belief that the manager is good following observation of the manager's first-period performance. Let P_{1B} (which equals $1 - P_{1G}$) equal the owner's Bayesian posterior belief that the manager is bad. The cost of replacing the manager with a new manager for the second period is R . The manager will be retained if, following observation of the manager's performance in the first period, the owner's expected utility in the second period with retention less the cost of replacement is greater than or equal to the expected utility of replacement, that is, if

$$(3.3a) \quad (P_{1G}) EU(Y|G) + (P_{1B}) EU(Y|B) \geq (P_G) EU(Y|G) + (P_B) EU(Y|B) + R.$$

Using $P_G + P_B = 1$ and $P_{1G} + P_{1B} = 1$,

$$(3.3b) \quad P_{1G} - P_G \geq R/(EU(Y|G) - EU(Y|B)).$$

The manager is retained when P_{1G} is large enough to make the left-hand side of (3.3b) larger than the right-hand side. Let θ_0 equal P_G/P_B , and let θ_1 equal P_{1G}/P_{1B} . Since P_{1G} is an increasing function of θ_1 , there is also a unique value for theta, θ^* , such that the manager should be retained if $\theta_1 \geq \theta^*$. If $R = 0$, then the decision rule simplifies to $P_{1G} - P_G \geq 0$ and the criterion for retention can only (and will always) be met as long as $\theta_1 \geq \theta_0$. As long as $R = 0$, the owner's decision rule must be to choose a region, C , for which $\theta_1 \geq \theta_0$ when y is outside of C . For convenience, we will assume for the rest of the paper that $R = 0$.¹

The critical region that maximizes the owner's expected utility given the manager's first-period choice of variance will be known as C^* . To identify C^* , the best reply critical region, recall that, by Bayes' rule,

$$(3.4) \quad \theta_1 / \theta_0 = \frac{l(y|\mu_G(\sigma_{G^*}))}{l(y|\mu_B(\sigma_{B^*}))}.$$

When $R = 0$, the values of Y for which the ratio of likelihood functions given manager choices of variance is greater than or equal to one are within C^* . The owner's best reply to σ_G^* and σ_B^* will be

$$C^* = \left\{ y \mid \theta_1 < \theta_0 \frac{l(y \mid \mu_G(\sigma_G^*))}{l(y \mid \mu_B(\sigma_B^*))} \right\}.$$

3.3. Equilibrium

A set of strategies $\{C^*, \sigma_G^*, \sigma_B^*\}$ will combine to form a Bayesian Nash equilibrium if, given the owner's choice of C^* , managers of either type will have a best response of σ_i^* and, given the managers' choices of σ_i^* and the owner's beliefs regarding the likelihood that the chosen manager is of a particular type, the owner will have a best response of C^* .²

For the case where $R = 0$, the set of values of y for which $\theta_1 < \theta_0$ is the critical region, C^* , of the test for manager retention. When $y \in C^*$, the owner will fire the manager. Given C^* , managers of each type will choose the level of variance, σ^* , that maximizes their chance of attaining a realization of Y outside of C^* . Given the managers' choices for variance, the ratio of likelihood functions, $l(y \mid \mu_G(\sigma_G^*)) / l(y \mid \mu_B(\sigma_B^*))$, must equal one at the boundaries of C^* . If these conditions are met, the strategies result in a Bayesian-perfect equilibrium.

Assuming that $R = 0$ and Y is normally distributed, we can exploit the fact that $l(y \mid \mu_G(\sigma_G^*)) / l(y \mid \mu_B(\sigma_B^*)) = 1$ at the boundaries of C^* in order to determine the nature of the critical region. Referring to Figures 3a and 3b, which depict the probability distribution functions of two managers, we see that the ratio of likelihood functions is equal to one at the points where the two functions intersect. In Figure 3a, $\sigma_G^* = \sigma_B^*$. When two normal distributions have the same level of variance and different means, there will be only one realization of Y for which the ratio of likelihood functions is equal to one. This value of Y will be half way between the two

distributions' means. In this case, C has a single boundary, y^* . Realizations of Y less than y^* result in replacement of the manager. All realizations of Y greater than or equal to y^* result in retention of the manager. The critical region C will be the open interval $[-\infty, y^*)$ while \bar{C} will be the open interval $[y^*, \infty]$. In this case, μ_G must be inside of \bar{C} while μ_B is inside of C .

If two normal distributions have different levels of variance, there will be two realizations of Y for which the ratio of likelihood functions equals one, regardless of the distributions' means. In Figure 3b, if $\sigma_G^* > \sigma_B^*$, then realizations of Y between the two points where the ratio equals one will be part of the critical region. Hence, C^* will be the open interval (y^*, y^{**}) . The bad manager's choice of mean, μ_B , must be inside of C . If $\sigma_G^* < \sigma_B^*$, then realizations of Y between the two points where the ratio equals one will not be part of the critical region and \bar{C}^* will be the closed interval $[y^*, y^{**}]$. In this case, μ_G must be inside of \bar{C} .

The special case in which $R = 0$ has some other interesting properties. When there are no replacement costs, neither the true distribution of managers nor the owner's utility function plays any role in determining equilibrium. The owner's decision rule is to retain the manager as long as $\theta_1 \geq \theta_0$. By Bayes' rule, this will be the case only (and always) at points where the ratio of likelihood functions is greater than or equal to one. The true distribution of managers is not an argument in the likelihood functions of either type of manager. Hence, the true distribution of managers cannot affect the managers' choices of σ^* or the owner's choices of y^* and y^{**} .

To see why this special case is unaffected by changes in the owner's utility function, remember that the owner, though concerned with expected utility, is only able to observe realized utility. The likelihood that a manager of a given type produces a particular level of realized utility is identical to the probability that the manager produces the value of Y associated with that level of utility. Hence, the values of Y

for which ratio of likelihood functions is greater than or equal to one are the same as the values of $u(y)$ for which the ratio is greater than or equal to one. As a result, the owner will choose the same values for y^* and y^{**} (and the managers will respond by choosing the same values for σ^*) regardless of the owner's utility function.

4. A CLASS OF TRACTABLE EXAMPLES

When Y is normally distributed and the mean/variance frontier is of the quadratic type shown in (4.1a) and (4.1b), the manager retention problem can be solved analytically and graphically. Let the mean/variance frontier of a good manager be

$$(4.1a) \quad \mu_G = b\sigma_G - \sigma_G^2 \quad b > 1,$$

and let the frontier of a bad manager be

$$(4.1b) \quad \mu_B = \sigma_B - \sigma_B^2.$$

Neither manager will be allowed to set both μ and σ equal to zero.³ Figure 4 shows these frontiers for the case $b = 2$.

Assume that the critical region C is an interval of the form $[-\infty, y^*]$. Later, it will be demonstrated that this is, indeed, true for this class of examples. Given his choice of mean and variance, a manager would have performance Y , normally distributed with mean μ and variance σ^2 . Letting $\Phi(\cdot)$ be the standard normal cumulative density function (CDF), the likelihood of performance y^* or less is given by $\Phi(z)$, where $z = (y^* - \mu)/\sigma$.

The manager will choose the value of σ that minimizes the probability of his being fired. The probability of being fired is Φ . Since Φ is a monotonic increasing transformation of z , choosing σ to minimize Φ is the same as choosing σ to minimize

z. Substituting the values for the managers' means given by (4.1a) and (4.1b), the managers' choice problems are given by (4.2a) and (4.2b):

$$(4.2a) \quad \sigma_G^* = \operatorname{argmin} (y^* - b\sigma_G + \sigma_G^2)/\sigma_G$$

w.r.t. σ_G

$$(4.2b) \quad \sigma_B^* = \operatorname{argmin} (y^* - \sigma_B + \sigma_B^2)/\sigma_B.$$

w.r.t. σ_B

Setting the derivatives of (4.2a) and (4.2b) equal to zero and solving for σ^* , we find that both types of managers' optimal choice is to set σ^* equal to $(y^*)^{1/2}$. As a result, $\mu_G^* = b(y^*)^{1/2} - y^*$, and $\mu_B^* = (y^*)^{1/2} - y^*$. Since both managers choose the same value for σ^* , there is, indeed, only one value of Y , y^* , where the ratio of likelihood functions is equal to one and the critical region is of the form $[-\infty, y^*)$.

The managers' decision-making process can readily be represented graphically in mean/variance space. In Figure 4, several of the managers' indifference curves are shown. Managers are interested in choosing the mean/variance combination that achieves the lowest value of z . Managers will be indifferent to all combinations of μ and σ that produce the same value of z leading to linear indifference curves of the following form:

$$(4.3) \quad \mu = -z\sigma + y^*.$$

All of the indifference curves converge at the point where $\mu = y^*$ and $\sigma = 0$. At this point, z is undefined. From this point, the indifference curves fan out. The indifference curve to the left of y^* along the μ axis represents the locus of mean/variance combinations that have a zero probability of surpassing y^* . The indifference curve starting at y^* and moving to the right along the μ axis represents the locus of

mean/variance combinations which assure that y^* will be equaled or surpassed. The indifference curve that rises vertically from y^* , labeled I_2 , represents all mean/variance combinations that have a 50 percent chance of surpassing y^* . Indifference curves grow more preferable as they fan from left to right, so a manager will choose the combination on his mean/variance frontier tangent to the indifference curve that is most to the right.

Given the managers' choices for σ^* , it is not difficult to identify the owner's choice of y^* . Recalling that we have assumed that replacement costs are equal to zero, we know from (3.4) that y^* must be at the point where $l_R/l_B = 1$. The derivative of Φ , ϕ , is the normal distribution's probability density function (PDF). The l_R will equal ϕ evaluated at $\mu_G^*(\sigma_G^*)$ and σ_G^* while l_B will equal ϕ evaluated at $\mu_B^*(\sigma_B^*)$ and σ_B^* . Setting the normal likelihood functions for both types of managers equal to one another, substituting $(y^*)^{1/2}$ for σ^* , and simplifying, we get the equation for y^* :

$$(4.4) \quad [b(y^*)^{1/2} - 2y^*]^2 = [(y^*)^{1/2} - 2y^*]^2.$$

Solving (4.4) for y^* , we find that $y^* = (b + 1)^2/16$ and $\sigma^* = (b + 1)/4$. Substituting this value of σ^* into (4.1a) and (4.1b), we find that the mean for a bad manager is $(2b - b^2 + 3)/16$ while the mean for a good manager is $(6b - b^2 + 7)/16$. Evaluated at $b = 2$, the probability that a good manager will be retained is .81 while the probability that a bad manager will be retained is .19.

Referring back to Figure 4, we see a graphic representation of the equilibrium. Each manager chooses the (μ, σ) combination where their mean/variance frontier is tangent to the best attainable indifference curve. For a bad manager, this is I_1 while it is I_3 for a good manager. Both types of managers have chosen the same level of variance, but good managers have a higher mean than bad managers. Good managers are on the efficient portion of their mean/variance frontiers while bad managers are on the inefficient portion of their mean/variance frontiers.

5. POTENTIAL EQUILIBRIA

When we assume that Y is normally distributed and $R = 0$, it is possible to identify classes of strategy sets consistent with the equilibrium developed in section 3. In this section, we demonstrate through a process of elimination that any strategy set consistent with equilibrium must fall into one of three broad groups of strategy set classes. One of these groups includes the equilibrium presented in section 4. To show that equilibria from the remaining groups are possible, we present an example from each.

5.1. Elimination of Strategy Sets Inconsistent with Equilibrium

In order to eliminate strategy sets inconsistent with equilibrium, we exploit four lemmas. To facilitate the explanation of these lemmas, recall that, for any given value of μ less than the maximum value attainable by the manager, there are two (μ, σ) combinations on the managers mean/variance frontier. One of these combinations is a low-variance choice on the efficient side of the frontier while the other is a high-variance choice on the inefficient side of the frontier. Since the manager may choose (μ, σ) combinations inside of the area bounded by the σ axis and the frontier, the manager can choose, for a given value of μ , any level of variance between the high- and low-variance values on the manager's frontier. For a given value of μ , all values of σ between and including these two values are said to be feasible. By the same reasoning, for a given value of σ , all values of μ between zero and the maximum attainable value of μ given σ are said to be feasible.

LEMMA 1: As σ decreases, the probability mass near μ increases, which increases the probability of y lying between y^* and y^{**} that bracket the mean. The Appendix includes a formal proof of this fact about normals.

LEMMA 2: For a given value of σ , the probability of falling between y^* and y^{**} is maximized when $\mu(\sigma) = \bar{\mu} \equiv (y^* + y^{**})/2$. $\partial\Omega/\partial\mu$ is positive for $\mu < \bar{\mu}$ and negative for $\mu > \bar{\mu}$. This is a property of the normal distribution. See Appendix.

In terms of the rejection region C , the lemma implies that, if C is the open interval (y^*, y^{**}) , then managers will choose a mean as far from $\bar{\mu}$ as possible. If \bar{C} is the closed interval $[y^*, y^{**}]$, then managers will choose a mean as close to $\bar{\mu}$ as possible.

It can easily be shown that Lemma 2 remains true as y^{**} approaches ∞ and/or y^* approaches $-\infty$. Proof of this has been omitted.

LEMMA 3: If $\mu_G = \mu_B$, then μ_G and μ_B are equal to $\bar{\mu}$. Again, this is a property of the normal distribution. See Appendix.

LEMMA 4: Consider the case in which $\sigma_G < \sigma_B$. It is a property of normals, as shown in the Appendix, that, as μ_B increases (μ_G constant), $\bar{\mu}$ will decrease.

Consider the case in which $\sigma_G < \sigma_B$. If $d\bar{\mu}/d\mu_B$ is negative for all values of μ_B between y^* and y^{**} , then Lemma 4 is true. To see why this is the case, consider a situation in which $\mu_G > \mu_B$ and $d\bar{\mu}/d\mu_B$ is negative. As μ_B increases and converges on μ_G , $\bar{\mu}$ must be decreasing in value, but Lemma 3 assures that, when $\mu_G = \mu_B$, $\bar{\mu}$ will be equal to them. Hence, as μ_B converges on μ_G from below, $\bar{\mu}$ converges on μ_G from above. This implies that μ_G is between $\bar{\mu}$ and μ_B when $\mu_G > \mu_B$. Similar reasoning assures that μ_G is between $\bar{\mu}$ and μ_B when $\mu_G < \mu_B$.

By exploiting the lemmas presented above, we can prove five theorems that demonstrate that most Bayesian combinations of good and bad manager strategies cannot form part of a strategy set consistent with a Bayesian-perfect equilibrium. Theorem 1 eliminates all possible strategy sets in which managers of either type choose mean/variance combinations that do not lie on their mean/variance frontiers.

THEOREM 1: In equilibrium, managers must choose mean/variance combinations that lie on their mean/variance frontiers.

PROOF: Manager choices of μ and σ that do not lie on the manager's mean/variance frontier may either be on the σ axis (where μ is equal to zero and σ is between the lowest and highest values possible when $\mu = 0$) or interior to the region bounded by the mean/variance frontier and the σ axis.

First, consider interior choices. Regardless of the nature of the equilibrium, either C or \bar{C} or both will be intervals. If \bar{C} (C) is a closed (open) interval then, by Lemma 2, managers will always choose the feasible value of μ , given his choice of σ , that is as close to (far from) $\bar{\mu}$ as possible. There are six possible cases:

1-2. If $\bar{\mu} \leq 0$ and \bar{C} (C) is a closed (open) interval, then the manager will choose the lowest (highest) feasible value of μ . This choice is on the σ axis (mean/variance frontier).

3-4. If $\bar{\mu}$ is greater than or equal to the highest feasible level of μ given the manager's choice of σ and \bar{C} (C) is a closed (open) interval, then the manager will choose the highest (lowest) feasible value of μ . This choice is on the mean/variance frontier (σ axis).

5. If $\bar{\mu}$ is between zero and the highest feasible level of μ given the manager's choice of σ and \bar{C} is a closed interval, then the manager will choose a mean of $\bar{\mu}$. By Lemma 1, the manager will then choose the lowest feasible value of σ given his choice of μ . This choice is on the efficient portion of the manager's mean/variance frontier.

6. If $\bar{\mu}$ is between zero and the highest feasible level of μ given the manager's choice of σ and C is the open interval then, by Lemma 2, the manager will choose a value of μ as far from $\bar{\mu}$ as possible. This choice is either on the mean/variance frontier or the σ axis.

Now consider manager choices of (μ, σ) combinations on the σ axis but not on the mean/variance frontier. There are four possible cases:

1. If $\bar{\mu}$ is greater than zero and \bar{C} is a closed interval, then by Lemma 2, the manager will choose a value of μ closer to $\bar{\mu}$ than zero. This choice cannot be on the σ axis.

2. If $\bar{\mu}$ is less than or equal to zero and \bar{C} is a closed interval, then the manager will set μ equal to zero. If, however, managers of both types choose the same mean then, by Lemma 3, their choice of mean is equal to $\bar{\mu}$. Since the manager, regardless of type, chooses zero as his mean, $\bar{\mu} = 0$ and the manager's chosen mean is inside of \bar{C} . By Lemma 1, the manager will then choose the lowest feasible value of σ . This choice is on the efficient portion of the mean/variance frontier.

3. If $\bar{\mu}$ is less than or equal to one-half of the maximum feasible value of μ given the manager's choice of σ and C is an open interval then, by Lemma 2, the managers will choose the highest feasible value of μ . This choice is on the manager's mean/variance frontier.

4. If $\bar{\mu}$ is greater than one-half of the maximum feasible value of μ given the manager's choice of σ and C is an open interval, then the manager will set a mean of zero. Since managers, regardless of type, choose zero as their mean, Lemma 3 assures that $\bar{\mu} = 0$ and the manager's chosen mean is inside of C . By Lemma 1, the manager will then choose the highest feasible value of σ consistent with a mean of zero. This choice is on the inefficient portion of the manager's mean/variance frontier.

As has been shown, no manager choice of a (μ, σ) combination not lying on their mean/variance frontier can be consistent with an equilibrium of the type outlined in section 3. \diamond

Theorem 1 proves that a manager must choose a (μ, σ) combination lying on his mean/variance frontier if his choice is to be part of an equilibrium strategy set.

Remaining combinations of manager strategies can be characterized by three properties: (1) is μ_G greater than, less than, or equal to μ_B ; (2) is σ_G greater than, less than, or equal to σ_B ; and (3) on what portion of the mean/variance frontier are the (μ, σ) combinations chosen by the managers. There are no less than 81 possible combinations of properties 1-3. These combinations are shown in Table I.

Most of the classes of strategy sets listed in Table I cannot form part of an equilibrium of the type described in section 3. A number of the combinations are ruled out by assumptions that we have made regarding the nature of the managers' mean/variance frontiers. For example, there can be no equilibrium in which both managers choose the (μ, σ) combination at their transition points [where $(d\mu)/(d\sigma) = 0$], while $\mu_G \leq \mu_B$ since the mean/variance frontiers of good managers strictly dominate those of bad managers, assuring that $\mu_G > \mu_B$ at the transition points. Combinations ruled out by assumptions that we have made regarding the structure of the mean/variance frontiers are denoted by the symbol "A" in Table I.

We will now prove four additional theorems that eliminate many of the potential equilibria remaining in Table I. Combinations eliminated by the various theorems are denoted by the letter "T" followed by a number identifying the theorem that was applied to disqualify it. The remaining combinations are denoted by the letter "P" for possible Bayesian-perfect equilibria.

THEOREM 2: There can be no equilibrium in which C is defined by a single value $-y^*$, and good (bad) managers choose combinations on the inefficient (efficient) portion of their mean/variance frontiers.

PROOF: If C is defined by a single value, μ_G will always be outside of C while μ_B is always within C . By Lemma 1, the level of variance consistent with μ_G that maximizes the probability of a realization outside of C for a good manager is the lowest available. This cannot be on the inefficient portion of the good managers'

mean/variance frontier. By the same reasoning, the level of variance consistent with μ_B that minimizes the probability of a realization inside of C for a bad manager is the highest available. This cannot be on the efficient side of a bad manager's mean/variance frontier. \diamond

THEOREM 3: There can be no equilibrium strategy set in which $\sigma_G < \sigma_B$ ($\sigma_G > \sigma_B$) and good (bad) managers choose combinations on the inefficient (efficient) portion of their mean/variance frontiers.

PROOF: If $\sigma_G < \sigma_B$, then μ_G is between y^* and y^{**} and is outside of C . By Lemma 1, good managers should choose the lowest value for σ consistent with μ_G . This choice cannot be on the inefficient portion of the good managers' mean/variance frontier.

If $\sigma_G > \sigma_B$, then μ_B is between y^* and y^{**} and is inside of C . By Lemma 1, bad managers should choose the highest value for σ consistent with μ_B . This choice cannot be on the efficient portion of the bad managers' mean/variance frontier. \diamond

THEOREM 4: There can be no equilibrium strategy set in which $\mu_G < \mu_B$.

PROOF: If $\mu_G < \mu_B$ and $\sigma_G < \sigma_B$, recall from section 3 that \bar{C} is the closed interval $[y^*, y^{**}]$ and $y^* < \mu_G < y^{**}$. If μ_B is greater than y^{**} then, by Lemma 2, the bad manager would prefer to lower his choice of mean. He is, of course, able to do so since managers can always choose any (μ, σ) combination within the area bounded by the σ axis and their mean/variance frontiers. If μ_B is less than or equal to y^{**} then, by Lemma 4, μ_G is between $\bar{\mu}$ and μ_B . By Lemma 2, bad managers should prefer to lower their choice of mean at least to μ_G from μ_B .

If $\mu_G < \mu_B$ and $\sigma_G > \sigma_B$, then C is the open interval (y^*, y^{**}) and $y^* < \mu_B < y^{**}$. If μ_G is less than y^* then, by Lemma 2, the bad manager would prefer to lower his choice of mean. If μ_G is greater than or equal to y^* then, by Lemma 4, μ_B is

between $\bar{\mu}$ and μ_G . By Lemma 2, bad managers should prefer to choose μ_G rather than μ_B . \diamond

THEOREM 5: There can be no equilibrium strategy set in which good managers are on the efficient portion of their mean/variance frontiers while bad managers are on the inefficient portion of their mean/variance frontiers, $\mu_G = \mu_B$, and $\sigma_G < \sigma_B$.

PROOF: Since $\sigma_G < \sigma_B$, both managers' mean will be inside of \bar{C} . By Lemma 1, bad managers will choose the lowest feasible level of variance given their choices of mean. This choice cannot be on the inefficient portion of their mean/variance frontiers. \diamond

Of the 81 possible combinations in Table I, 24 of them remain as possible Bayesian-perfect equilibria. On inspection, these combinations fall into three categories. These categories are:

(A) $\sigma_G < \sigma_B$, $\mu_G \geq \mu_B$, no managers on the inefficient portion of their mean/variance frontiers.

(B) $\sigma_G > \sigma_B$, $\mu_G \geq \mu_B$, no bad manager on the efficient portion of his mean/variance frontier.

(C) $\sigma_G \leq \sigma_B$, $\mu_G \geq \mu_B$, no bad (good) managers on the efficient (inefficient) portion of their mean/variance frontiers.

The equilibrium presented in section 4 is an example of a type (C) equilibrium. Recall that, in the example of section 4, both managers set the same level of variance and $\mu_G > \mu_B$. The good manager's (μ, σ) choice is on the efficient portion of his mean/variance frontier while the bad manager's choice is on the inefficient portion of his frontier.

Generating examples of type (A) and type (B) equilibria, in which $\sigma_G \neq \sigma_B$, is difficult to do algebraically. Using a computer algorithm, we are able to demonstrate

the existence of equilibrium strategy sets of these types. These examples are presented graphically in Figures 6a and 6b.

As an example of a type (A) equilibrium, consider a situation in which the mean/variance frontiers of our managers are defined by

$$(5.1a) \quad \mu_G = 2(\sigma_G - \sigma_G^2)$$

$$(5.1b) \quad \mu_B = \sigma_B - \sigma_B^2.$$

The owner sets $\bar{C} = (.114, .483)$. Given their mean/variance frontiers, the choices of variance that will minimize the probability of a realization of Y inside of C are to set $\sigma_G = .128$ and $\sigma_B = .181$. As a result of these choices, $\mu_G = .223$ and $\mu_B = .148$. The probability that a good manager gets rehired is .782 while the probability of a bad manager getting rehired is .543. The ratio of likelihood functions, L , for the distributions chosen by the managers is equal to one at the points .114 and .484, completing the equilibrium. The equilibrium is, indeed, of type (A) since both managers' (μ, σ) choices are on the efficient portion of their mean/variance frontiers, the good manager has chosen a lower level of variance than the bad manager, and $\mu_G > \mu_B$.

As an example of a type (B) equilibrium, consider a situation in which the mean/variance frontiers of our managers are defined by (4.1a) and (4.1b). The owner chooses $C = (-1.36, 1.36)$. Both managers maximize their probability of being outside of C by choosing their maximum attainable variance – $\sigma_G = 2$ and $\sigma_B = 1$. Both managers' mean is zero. The probability that a good manager gets rehired is .496 while the probability of a bad manager getting rehired is .174. The likelihood ratio L is equal to one at the points -1.36 and 1.36 , completing the equilibrium. The equilibrium is, indeed, of type (B) since both managers' (μ, σ) choices are on the inefficient

portion of their mean/variance frontiers, the good manager has chosen a higher level of variance than the bad manager, and $\mu_G = \mu_B$.

6. SOME WELFARE IMPLICATIONS OF THE EQUILIBRIUM CONCEPT

It is possible to demonstrate that, assuming the owner's objective is to maximize expected Y (EY) over both periods, an increase in the quality of good managers has an *ambiguous* effect on the owner's welfare while an increase in P_G (the proportion of good managers in the population) unambiguously enhances the owner's welfare.

THEOREM 6: An increase in P_G (the proportion of the manager population that is good) increases EY over both periods.

PROOF: To prove Theorem 6, we must demonstrate that $d(EY)/dP_G$ is positive. Let π_G (π_B) equal the probability that a good (bad) manager chosen in the first period will be retained. Let μ_{G2} (μ_{B2}) be the mean level of Y produced by a good (bad) manager in the second period. Let $A = P_G(\mu_{G2}) + (1 - P_G)(\mu_{B2})$, the expected return of a manager chosen at random during the second period. Then, (6.1) is the equation for EY over both periods:

$$(6.1) \quad EY = P_G \cdot \mu_G + P_G \pi_G \cdot \mu_{G2} + (1 - \pi_G) P_G \cdot A + (1 - P_G) \mu_B \\ + (1 - P_G) \pi_B \cdot \mu_{B2} + (1 - P_G) (1 - \pi_B) \cdot A.$$

Recall from section 3 that the distribution of managers has no effect on equilibrium when replacement costs equal zero. Hence, $d\pi_G/dP_G$ and $d\pi_B/dP_B$ are equal to zero. As a result, the derivative of (6.1) with respect to P_G is

$$(6.2) \quad d(EY)/dP_G = \mu_G + \pi_G \cdot \mu_{G2} + (1 - \pi_G) [P_G(dA/dP_G) + A] \\ - \mu_B - \pi_B \cdot \mu_{B2} + (1 - \pi_B) [(1 - P_G)(dA/dP_G) - A].$$

To show that the sign of (6.2) is positive, the terms of the equation can be rearranged and simplified:

$$\begin{aligned}
 (6.3) \quad d(EY)/dP_G &= [\mu_G - \mu_B] + [(1 - \pi_G) P_G \cdot (dA/dP_G)] \\
 &+ [(1 - \pi_B) (1 - P_G) (dA/dP_G)] + [(\pi_G \cdot \mu_{G2}) \\
 &- \pi_G A] + [-(\pi_B \cdot \mu_{B2}) + \pi_B A].
 \end{aligned}$$

The first term is positive, since Theorem 4 establishes that $\mu_G > \mu_B$. The second and third terms are positive, since dA/dP_G is equal to $\mu_{G2} - \mu_{B2}$ and $\mu_{G2} > \mu_{B2}$ by assumption. Since A is just a weighted average of μ_{G2} and μ_{B2} , $\mu_{G2} > A > \mu_{B2}$. As a result, the last two expressions are positive and $d(EY)/dP_G$ is positive. \diamond

Although $d(EY)/dP_G$ is always positive, there are circumstances in which a risk-averse owner may be worse off if the proportion of good managers increases. For example, consider our example of a type (B) equilibrium and a risk-averse owner with the quadratic utility function $U = -(1 - Y)^2$. Since the owner's utility function has no effect on first-period manager choices when there are no replacement costs, managers make variance choices identical to those in our example – $\sigma_G = 2$ and $\sigma_B = 1$. For both types of manager, mean return equals zero. As a result, the owner's first-period expected utility from good managers is -5 while it is -2 for bad managers. In this case, the good manager has identified himself by performing substantially worse than the bad manager. Since managers of either type choose to maximize the owner's expected utility during the second period, good managers will set σ approximately to .4 (resulting in an expected utility of $-.16$) while bad managers set σ approximately to .33 (resulting in an expected utility of $-.667$). Substituting these values into (5.1) in place of the first- and second-period mean returns, and using the probabilities of retention for both manager types provided by the type (B) example, we find that $EU = -2 - 4.28P_G + .8(P_G)^2$. The derivative of this expression w.r.t. P_G is $-4.28 + 1.6P_G$, which is negative for any conceivable value of P_G . Hence, in this example, an increase

in the proportion of good managers lowers the owner's expected utility over both periods.

THEOREM 7: An increase in the ability of good managers has an ambiguous effect on EY over both periods.

PROOF: In order to prove Theorem 7, we provide an example in which an increase in the ability of good managers results in either an increase or a decrease in EY over both periods depending on the distribution of managers in the population.

Consider the example presented in section 4. In that example, the managers' mean/variance frontiers were defined as $\mu_G = b\sigma_G - \sigma_G^2$ and $\mu_B = \sigma_B - \sigma_B^2$. Assume that $b = 2$. Managers of both types will make the choices outlined above in the example of section 4. By incorporating into (6.1) the resulting values for first- and second-period mean returns for both types of manager as well as their probabilities of survival, we find that

$$(6.4) \quad EY = 1.53P_G + .44.$$

Now let us assume that there is an increase in the ability of good managers so that $b = 3$. By incorporating the resulting values for first- and second-period returns and probabilities of survival into (6.1), we find that

$$(6.5) \quad EY = 4.17P_G + .08.$$

Comparing (6.4) and (6.5), we see that EY's relationship to the value of b depends critically on the value of P_G . If $P_G < .125$, then an increase in the ability of good managers from $b = 2$ to $b = 3$ lowers EY over both periods. If $P_G > .125$, then an increase in the ability of good managers from $b = 2$ to $b = 3$ raises EY over both periods. \diamond

The intuition behind Theorem 7 is that increases in the ability of good managers force bad managers into taking ever more desperate gambles. This can offset all the other benefits of the increase in ability. When the value of b is raised from two to three, the managers' choice of σ increases from $3/4$ to one. As a result, μ_B declines from $3/16$ to zero while μ_G increases from $15/16$ to two. The expected second-period return for good managers, μ_{G2} , increases from one to 2.25 while μ_{B2} is unchanged. The probability that a good manager will be retained following observation of first-period results increases from .69 to .84 while the probability of bad managers retention falls from .31 to .16. If bad managers are far more common than good managers, the first-period decline in mean return for them far outweighs the first-period gain in mean return for good managers. Since good managers are rare, increases in their second-period mean return or in the probability that good (bad) managers are retained (replaced) hardly matter. Replacing a bad manager is meaningless if his replacement is almost inevitably just as bad a manager.

7. CONCLUSIONS

We have demonstrated that conditions may exist that induce relatively capable managers, influenced by reputational concerns, to behave in a manner that owners would regard as overly cautious while the same conditions and concerns induce less capable managers to behave in a manner that owners would regard as overly aggressive. We call equilibria in which such behavior is seen endgame equilibria. Such behavior is likely to take place in a wide variety of real-world managerial settings and can easily be induced when managers unobservably choose investments or business strategies that, combined with their level of ability, generate a stochastic stream of profits, whose realizations are then used to draw inferences regarding manager ability.

Given that the owners' interest is in the maximization of profits, the equilibrium behavior described in section 3 is costly. Bad managers choose a greater-than-profit-maximizing level of risk, sacrificing mean profits in order to increase the standard deviation of profits, while good managers choose a lower-than-profit-maximizing level of risk. Furthermore, the equilibrium behavior of managers would not be affected if, rather than profit maximization, owners were interested in maximizing some reasonable utility function. This is because the point where the ratio of likelihood functions equals one would remain unchanged regardless of the owners' utility function. As long as managers choose the same value for σ^* , the owners' best response, regardless of their utility functions, remains the same.

In this paper, we severely constrained the contract choices that owners could offer managers. The only type of contract possible lasted a single period and involved a fixed payment. This allowed us to greatly simplify the model's presentation and permitted us to focus on the incongruity between owner and manager preferences. The design of appropriate contractual mechanisms aimed at mitigating endgame-type problems, as well as the influence of such behavior on the firms' decisions regarding capital structure, remains for future research.

In considering the design of contractual mechanisms aimed at mitigating endgame behavior, it should be appreciated that, unlike standard contingent contracting models, owners attempting to mitigate this type of agency effect may prefer managers to be more, rather than less, risk averse. The reason for this is obvious. In typical contingent contracting, the efficacy of the contract is limited by the degree of manager risk aversion. Owners, assumed to be risk neutral, always wish that managers would be less risk averse and, hence, more willing to accept a share of a stochastic stream of profits. After all, the greater the degree to which manager compensation is tied to profits, the better the alignment of manager and owner preferences.

In the endgame model, however, bad managers behave too aggressively. These managers may not be so willing to sacrifice mean profits in order to increase the variance of profits if they are sharing in those profits. The manager's desire to raise mean income will induce more cautious behavior while the desire to avoid risks will further increase the degree of caution chosen in making decisions regarding the mean and variance of profits. This is exactly what owners want. Hence, depending on the frequency of good and bad managers, owners may see manager risk aversion as facilitating, rather than hindering, the design of contractual mechanisms which will reduce the incongruity of preferences between owners and managers.

FOOTNOTES

¹If both agents choose the same combinations of mean and variance, the equilibrium breaks down.

²For a discussion of Bayesian Nash equilibria, see Gibbons (1992).

³This assumption is required in order to assure that the mean/variance frontiers of good agents strictly dominate bad agents.

REFERENCES

- Chevalier, J., and G. Ellison (1995): "Risk Taking by Mutual Funds as a Response to Incentives," Working Paper, University of Chicago.
- Falkenstein, E.G. (1996): "Preferences for Stock Characteristics as Revealed by Open-End Mutual Fund Portfolio Holdings," *Journal of Finance*, 51, 111-135.
- Fama, E. (1980): "Agency Problems and the Theory of the Firm," *Journal of Political Economy*, 88, 288-307.
- Gibbons, R. (1992): *Game Theory for Applied Economists*. Princeton: Princeton University Press.
- Gibbons, R., and K.J. Murphy (1992): "Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence," *Journal of Political Economy*, 100, 468-505.
- Holmstrom, B. (1982): "Managerial Incentive Problems: A Dynamic Perspective," in *Essays in Economics and Management in Honor of Lars Wahlbeck*. Helsinki: Swedish School of Economics.
- Holmstrom, B., and J. Ricart i Costa (1986): "Managerial Incentives and Capital Management," *Quarterly Journal of Economics*, 101, 835-860.
- Huddart, S. (1996): "Reputation and Performance Fee Effects on Portfolio Choices by Investment Advisors," *Journal of Financial Abstracts: Series A, Corporate Finance and Organization Working Paper Series*, 3.
- Ricart i Costa, J.E. (1989): "On Managerial Contracting with Asymmetric Information," *European Economic Review*, 33, 1805-1829.
- Scharfstein, D.S., and J.C. Stein (1990): "Herd Behavior and Investment," *American Economic Review*, 80, 465-479.
- Zwiebel, J. (1995): "Corporate Conservation and Relative Compensation," *Journal of Political Economy*, 103, 1-25.

APPENDIX

PROOF OF LEMMA 1: Since Y is normally distributed, the probability of a realization of Y between y^* and y^{**} is given by

$$(A.1) \quad \Omega = \Phi[(y^{**} - \mu)/\sigma] - \Phi[(y^* - \mu)/\sigma],$$

where Φ is the standard normal CDF and is a function of the z values associated with the standard normal distribution. The derivative of Ω w.r.t. σ is

$$(A.2) \quad \partial\Omega/\partial\sigma = (\phi|_{y^{**}}) (\mu - y^{**})/\sigma^2 + (\phi|_{y^*}) (y^* - \mu)/\sigma^2,$$

where ϕ is the standard normal PDF evaluated at either y^* or y^{**} . Since ϕ and σ are always positive, expression (A.2) is negative when $y^* \leq \mu \leq y^{**}$, so the probability of falling between y^* and y^{**} decreases as σ increases when $y^* \leq \mu \leq y^{**}$. If $y^* \leq \mu \leq y^{**}$, $\partial\Omega/\partial\sigma$ is negative, and a manager should choose the lowest (highest) feasible value of σ should he seek to maximize (minimize) the probability of a realization of Y between y^* and y^{**} .

Recall from section 3 that, if $\sigma_B < \sigma_G$, then C is the open interval (y^*, y^{**}) . In this case, managers whose choice of mean is between y^* and y^{**} will choose the highest possible value for σ . If $\sigma_B > \sigma_G$, then \bar{C} is the closed interval $[y^*, y^{**}]$ and managers whose choice of mean is between y^* and y^{**} will choose the lowest possible value for σ . If $\sigma_B = \sigma_G$, then both C and \bar{C} are open intervals and managers whose choice of mean is inside of C (\bar{C}) will choose the highest (lowest) possible value for σ .

It can easily be shown that Lemma 1 remains true as y^{**} approaches ∞ and/or y^* approaches $-\infty$. Proof of this has been omitted. \diamond

PROOF OF LEMMA 2: The derivative of Ω w.r.t. μ is

$$(A.3) \quad \partial\Omega/\partial\mu = (\phi[(\mu - y^{**})/\sigma]) (-1/\sigma) + (\phi[(\mu - y^*)/\sigma]) (1/\sigma).$$

This expression is only equal to zero at $\bar{\mu} = (y^* + y^{**})/2$, since that is the only value of μ for which $\phi|_{y^{**}} = \phi|_{y^*}$. The second derivative of Ω w.r.t. μ is

$$(A.4) \quad [\phi|_{y^{**}}] (\mu - y^{**})/\sigma^3 - [\phi|_{y^*}] (\mu - y^*)/\sigma^3.$$

This expression is negative when evaluated at $\mu = (y^* + y^{**})/2$, so $\bar{\mu}$ is a local optimum. Since there is no other point where $\partial\Omega/\partial\mu = 0$, it is also a global optimum, which establishes the lemma. \diamond

PROOF OF LEMMA 3: Recall from section 3 that, when $\sigma_G \neq \sigma_B$, there are two points, y^* and y^{**} , where $l(y|\mu_G(\sigma_G^*))/l(y|\mu_B(\sigma_B^*)) = 1$. Because normal distributions are symmetric, y^* and y^{**} must be equidistant from $\mu_G = \mu_B$ and $y^* < \mu_G = \mu_B < y^{**}$. Hence, $\bar{\mu} = (y^* + y^{**})/2$ must equal μ_G and μ_B . Lemma 3 is illustrated graphically in Figure 5. \diamond

PROOF OF LEMMA 4: Since $\bar{\mu} = (y^* + y^{**})/2$, $d\bar{\mu}/d\mu_B$ will be negative if $dy^*/d\mu_B$ and $dy^{**}/d\mu_B$ are negative. First, consider $dy^*/d\mu_B$. Recall that in equilibrium

$$(A.5) \quad l(y^*|\mu_G(\sigma_G^*), \sigma_G^*) - l(y^*|\mu_B(\sigma_B^*), \sigma_B^*) = 0.$$

Taking the total derivative of (5.5) with respect to y^* and μ_B ,

$$(A.6) \quad [dl(y^*|\mu_G)/dy^* - dl(y^*|\mu_B)/dy^*] \Delta y^* = [dl(y^*|\mu_B)/d\mu_B] \Delta \mu_B.$$

Since $\mu_B > y^*$, $dl(y^*|\mu_B)/d\mu_B$ is negative. Referring to Figure 5, we see that while the slopes of both managers' probability distribution functions are positive at y^* , the slope is steeper for good managers. Hence, $dl(y^*|\mu_G)/dy^* > dl(y^*|\mu_B)/dy^*$ and

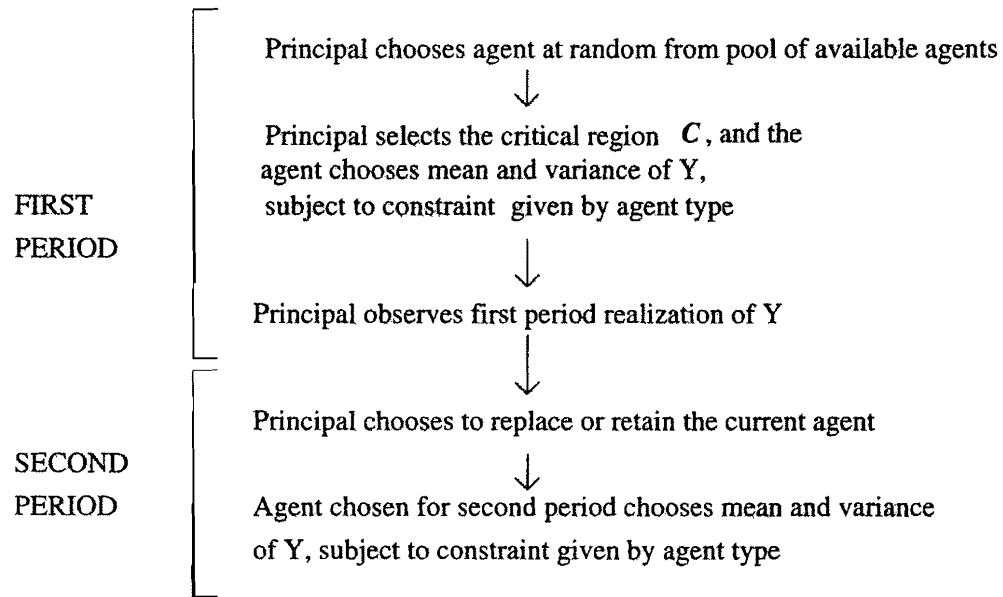
$dy^*/d\mu_B$ is negative. Following similar reasoning, $dI(y^{**}|\mu_B)/d\mu_B$ is positive and, while the slopes of both managers' probability distribution functions are negative at y^{**} , the slope is steeper for good managers. Hence, $dI(y^{**}|\mu_G)/dy^{**} < dI(y^{**}|\mu_B)/dy^{**}$ and both terms are negative. As a result, $dy^{**}/d\mu_B$ is negative. Since both $dy^*/d\mu_B$ and $dy^{**}/d\mu_B$ are negative, $d\bar{\mu}/d\mu_B$ is negative. Hence, μ_G is between $\bar{\mu}$ and μ_B . The proof for the case in which $\sigma_G > \sigma_B$ is similar and will not be shown. \diamond

TABLE I
POSSIBLE EQUILIBRIA

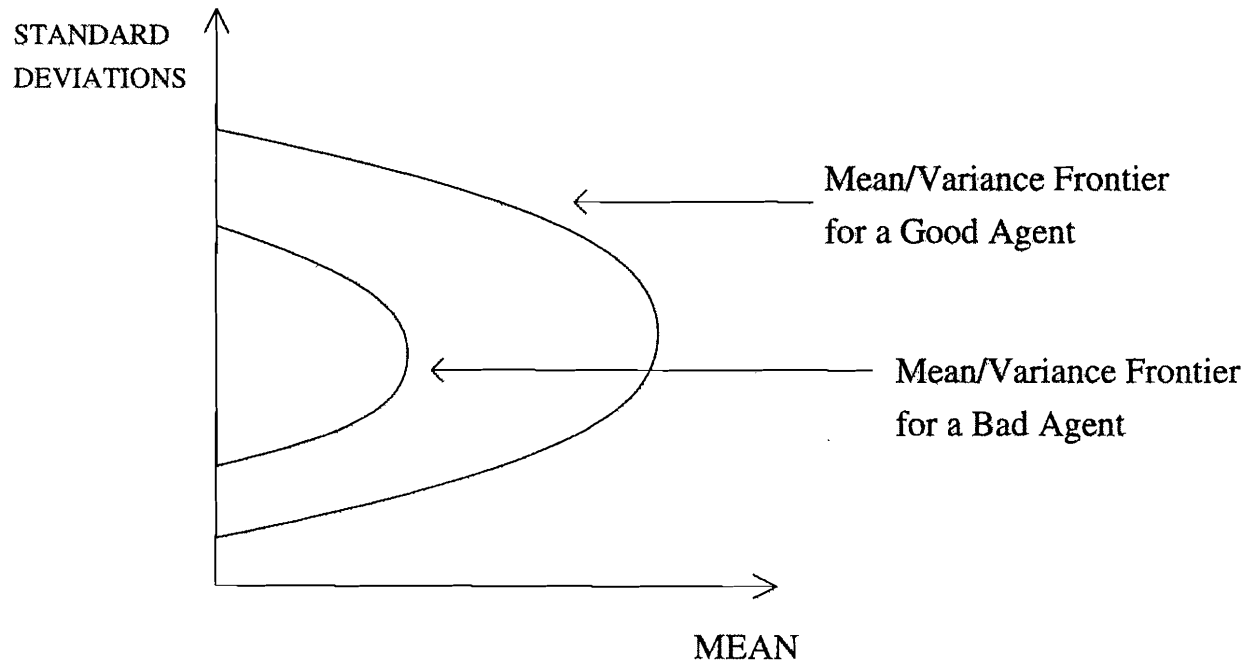
		$\mu_G > \mu_B$	$\mu_G = \mu_B$	$\mu_G < \mu_B$
G and B on inefficient side	$\sigma_G > \sigma_B$	possible	possible	T4
	$\sigma_G = \sigma_B$	T2	A	A
	$\sigma_G < \sigma_B$	T3	T3	T3
G on efficient side and B on transition	$\sigma_G > \sigma_B$	possible	possible	A
	$\sigma_G = \sigma_B$	possible	possible	A
	$\sigma_G < \sigma_B$	possible	possible	T4
G on inefficient side and B on efficient side	$\sigma_G > \sigma_B$	T3	T3	T3
	$\sigma_G = \sigma_B$	T2	A	A
	$\sigma_G < \sigma_B$	T3	T3	T3
B on inefficient side and G on transition	$\sigma_G > \sigma_B$	possible	A	A
	$\sigma_G = \sigma_B$	possible	A	A
	$\sigma_G < \sigma_B$	possible	A	A
B on inefficient side and G on efficient side	$\sigma_G > \sigma_B$	possible	A	A
	$\sigma_G = \sigma_B$	possible	A	A
	$\sigma_G < \sigma_B$	possible	T5	T4
G and B on efficient side	$\sigma_G > \sigma_B$	T3	T3	T3
	$\sigma_G = \sigma_B$	T2	A	A
	$\sigma_G < \sigma_B$	possible	possible	T4
B on efficient side and G on transition	$\sigma_G > \sigma_B$	T3	A	A
	$\sigma_G = \sigma_B$	T2	A	A
	$\sigma_G < \sigma_B$	possible	A	A
G and B on transition	$\sigma_G > \sigma_B$	possible	A	A
	$\sigma_G = \sigma_B$	possible	A	A
	$\sigma_G < \sigma_B$	possible	A	A

GRAPH 1: THE FLOW OF EVENTS

35



GRAPH 2: EXAMPLES OF MEAN/VARIANCE FRONTIERS



36

GRAPHS 3A AND 3B: DETERMINATION OF y^* AND y^{**}

37




Figure 3a: Sigma is the same for good and bad agents

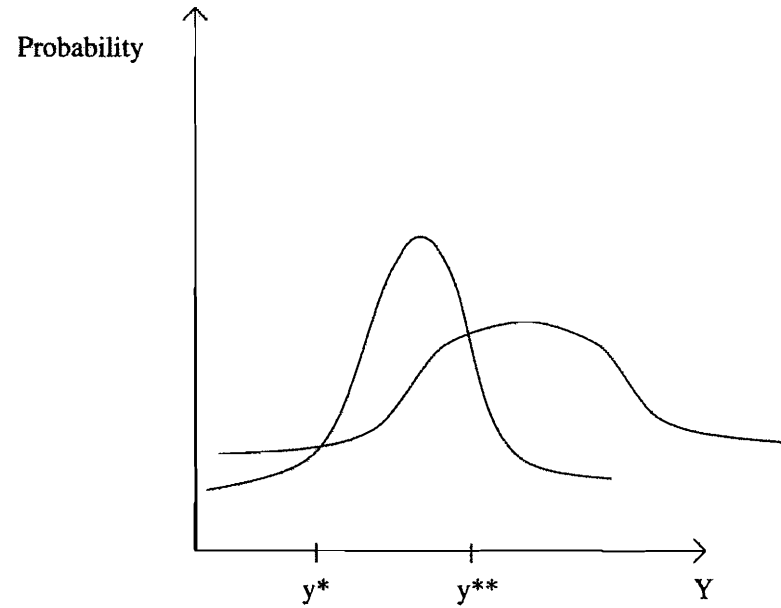




Figure 3b: Sigma is greater for one type of agent

GRAPH 4: A GRAPHIC EXAMPLE OF EQUILIBRIUM BEHAVIOR



38

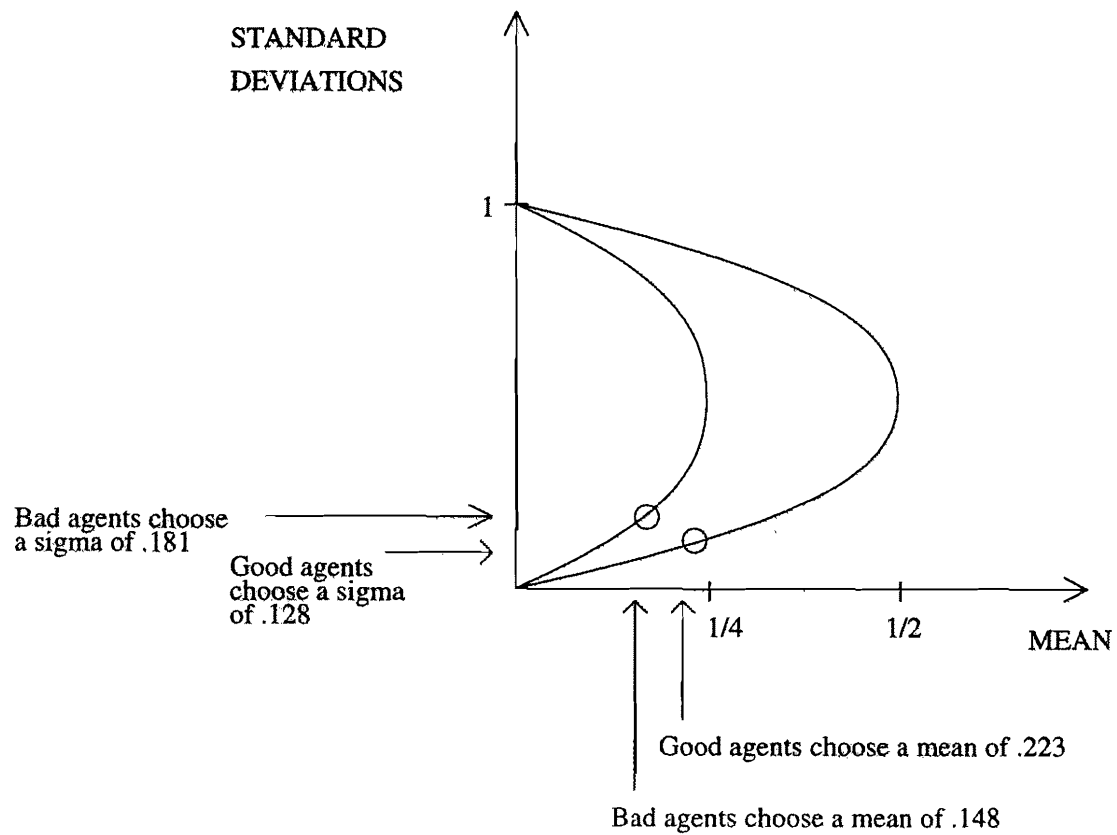
GRAPH 5: y^* , y^{**} , AND THE AGENTS' CHOICES OF MEAN



Both good and bad agents choose the same mean. Both y^* and y^{**} are equidistant from this mean.

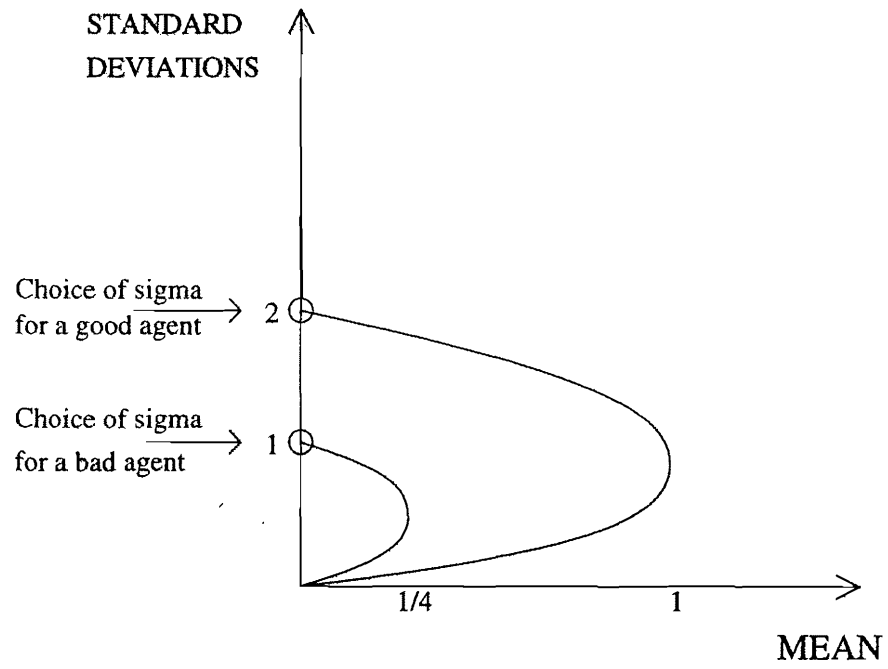
69

GRAPH 6A: A TYPE "A" EQUILIBRIUM



117

GRAPH 6B: A TYPE "B" EQUILIBRIUM



111