THE JOURNAL OF FINANCE • VOL. LXIII, NO. 6 • DECEMBER 2008

Overconfidence, CEO Selection, and Corporate Governance

ANAND M. GOEL and ANJAN V. THAKOR*

ABSTRACT

We develop a model that shows that an overconfident manager, who sometimes makes value-destroying investments, has a higher likelihood than a rational manager of being deliberately promoted to CEO under *value-maximizing* corporate governance. Moreover, a risk-averse CEO's overconfidence enhances firm value up to a point, but the effect is *nonmonotonic* and differs from that of lower risk aversion. Overconfident CEOs also underinvest in information production. The board fires both excessively diffident and excessively overconfident CEOs. Finally, Sarbanes-Oxley is predicted to improve the precision of information provided to investors, but to reduce project investment.

As you go the way of life you will see a great chasm. Jump. It is not as wide as you think. *Native American proverb.*

CHIEF EXECUTIVE OFFICERS (CEOs) AFFECT the quality of the information available to the board of directors and investors (e.g., Adams and Ferreira (2007) and Song and Thakor (2006)) as well as corporate investment decisions. Their personal attributes and behavioral biases, such as overconfidence, affect both their information-provision incentives as well as their investment decisions (e.g., Malmendier and Tate (2005)).

The interaction between the information-provision incentives of CEOs and the effectiveness of corporate governance is well recognized, both in the Sarbanes-Oxley Act (Sarbox) as well as in the recent corporate governance research (e.g., Adams and Ferreira (2007), Almazan and Suarez (2003), Harris and Raviv (2008), Hermalin and Weisbach (1998, 2003), Lorsch and MacIver (1989), Shleifer and Vishny (1997), and Song and Thakor (2006)). While this recent literature has yielded numerous valuable insights, it continues to leave unanswered puzzling questions like: Why do overconfident agents become CEOs despite the documented behavioral distortions associated with overconfidence (see, for example, Malmendier and Tate's (2005) evidence on overinvestment by overconfident CEOs)?

*Anand Goel is at DePaul University and Anjan Thakor is at Washington University, St. Louis. The helpful comments of Arnoud Boot; Fenghua Song; Josh Coval; an associate editor; an anonymous referee; the editor Rob Stambaugh; and Simon Gervais, the discussant of the paper at the American Finance Association's 2007 annual meeting in Chicago, are gratefully acknowledged.

We believe such puzzles exist because we do not have an analysis of the behavioral biases of those who are in the pool of managers from which CEOs are chosen, or of how corporate governance affects the composition of this pool. When a board of directors ("board" henceforth) is deciding who to appoint as CEO—whether from within the firm or from outside—attention typically focuses on a small set of senior executives who have survived the internal governance process by which managers get promoted through the corporate hierarchy. If this process is biased in favor of those with a specific set of attributes, then CEOs will display a preponderance of these attributes, regardless of the specific selection mechanism adopted by the board. And because CEO attributes affect various corporate decisions, the CEO *selection process* can affect the firm's investment policy as well as the efficacy of any corporate governance mechanism. It is therefore important to understand how this selection process works, how it is affected by initiatives like Sarbox, which managerial attributes and behavioral biases it tends to favor, and what this implies for corporate decisions.

We examine these issues by addressing the following questions: (1) When managerial ability is unknown and being inferred through observed performance, how does a constrained-efficient process of internal promotions, designed to select a CEO so as to maximize shareholder wealth, work? (2) How do managerial attributes and behavioral biases affect the probability of being promoted and can this explain why overconfident managers become CEOs? (3) How do the attributes of those who are endogenously promoted to be CEOs affect their firms' investment policies? (4) How do these attributes affect firm values? (5) How do these attributes affect CEOs' decisions about how much project-relevant information to acquire before deciding whether to invest? (6) If managers vying to be CEO vary cross-sectionally both in their risk aversion and their degree of overconfidence, is it possible to distinguish between overconfidence and reduced risk aversion? (7) How does the recognition of CEO attributes affect corporate governance at the board level, in terms of the board's decision to retain/fire the CEO contingent upon observed performance? (8) How does Sarbox affect the board's decision as to who to appoint as CEO and what does this imply for information disclosure, investment policy, and firm value?

In addressing these questions, we distinguish between "internal organizational governance" and "board governance." We view the former as referring to the internal promotion process by which managers move up through the corporate hierarchy and enter the pool from which one is eventually chosen to become CEO. We view the latter as the board's decision to retain or fire a CEO based on her observed performance. Both concepts of corporate governance have been studied separately in the literature and a significant focus of corporate governance reform is on the information, expertise, and incentives of the board to efficiently decide whether to retain or replace the CEO. However, our analysis reveals that the two processes are intertwined and that the board's assessment of the CEO must take into account the likely characteristics of the CEO as a result of the internal promotion process.

We develop a two-period leadership selection model in which there is initially one CEO, with many managers reporting to her, all risk averse.

Risk-neutral shareholders own the firm and there is a board of directors, acting in the shareholders' best interest, that watches over the CEO. In each period, the CEO chooses a "strategy" for the firm that affects the payoff distributions of all projects. Each manager chooses a project whose payoff distribution is affected by a choice made by the manager, the manager's ability, and the CEO's strategy. Each manager's ability is a priori unknown to all and is inferred over time from the observed payoffs of individual projects. The incumbent CEO retires at the end of the first period and the board replaces her with one of the managers who reported to her in the first period.

This basic model generates numerous results. First, the rational ability-filtering process associated with internal organizational governance generates an intrafirm tournament in which it is value-maximizing for the shareholders to appoint, as second-period CEO, the manager with the highest perceived ability at the end of the first period; this tournament induces each manager to take more first-period project risk than he would in the absence of the tournament.

Second, when an overconfident manager, that is, one who underestimates project risk, is introduced,[1] we find that the overconfident manager has the highest probability of being promoted to CEO when he is competing with otherwise rational managers.[2] Thus, the analysis implies that overconfidence is likely to be a more prevalent attribute among CEOs than in the general population. This result provides a possible explanation for why firms end up with overconfident CEOs even though such CEOs make value-destroying investments and CEO overconfidence is an *empirically detectable* attribute.[3]

We next examine how CEO overconfidence affects firm value. To do so, we consider a setting in which the CEO determines whether to invest in a portfolio of projects based on her private information about the portfolio payoff. We solve for the optimal compensation contract that trades off incentives for appropriate investment by the CEO against the cost of imposing risk on the risk-averse CEO.

Our third main result is that, under the optimal CEO compensation contract, a rational, risk-averse CEO underinvests in projects relative to the shareholders' optimum. This underinvestment reduces firm value. We show that a moderately overconfident risk-averse CEO increases firm value by mitigating the underinvestment problem. The reason is that an overconfident CEO overestimates the precision of her private information and overreacts to it. Thus, she invests in a project even when her positive information about the project is such that she would not invest in the project if she were rational.

[1] The finding that agents are overconfident is one of the most consistent in the psychology of judgment, (e.g., DeBondt and Thaler (1995)). A possible explanation for its persistence is offered by Bernardo and Welch (2001).

[2] We also consider a scenario with ex ante identical managers, where it is common knowledge that each manager may be overconfident with some probability. Taking this into account, the project risk each manager believes he has chosen is identical across all managers, but the true project risk is higher for overconfident managers. Consequently, an overconfident manager is more likely to outperform others and get promoted to CEO.

[3] In their analysis, Malmendier and Tate (2005) use publicly available data to detect CEO overconfidence. Presumably such detection is also within the reach of boards and investors.

Fourth, CEO overconfidence affects shareholder wealth *nonmonotonically* when CEOs are risk averse. While moderate overconfidence diminishes underinvestment and increases firm value, sufficiently high overconfidence generates overinvestment and decreases firm value. In contrast, firm value decreases *monotonically* in CEO risk aversion, which distinguishes overconfidence from lower risk aversion. The best outcome for the shareholders is thus to have a CEO who is overconfident but not too overconfident.

Fifth, an overconfident CEO *underinvests* in acquiring project-relevant information. This increases project selection errors and diminishes the quality of the information used to judge the CEO.

Sixth, we permit unobservable cross-sectional heterogeneity in both risk aversion and overconfidence among managers vying to be CEO, and design a truth-telling mechanism using the Revelation Principle that permits separation of managers by risk aversion but not overconfidence. This offers another perspective on the difference between overconfidence and risk aversion, and is in contrast to the existing literature which suggests that preferences and beliefs cannot be disentangled. For example, Sandroni and Squintani (2007) argue that overconfidence precludes the separation of low-risk and high-risk individuals by contract choice in insurance markets.

We then extend the model to focus on an incumbent CEO who is not destined to step down at the end of the first period but whose performance is being judged by the board to determine whether to retain her in the second period or replace her. Our seventh main result is that in a situation in which a project quality signal can be observed by the board in addition to the first-period project payoff, the board's decision to retain/fire the CEO depends on the interaction between its perception of the CEO's ability and its perception of the CEO's overconfidence. The board fires *both* low ability CEOs and those perceived to be either excessively cautious (diffident) or excessively overconfident.

Finally, we examine the impact of Sarbox on CEO selection, information provision incentives, investment policy, and firm value. We show that ex post penalties for providing imprecise information to investors can cause the board to shift its preference away from an overconfident manager when selecting a CEO. Thus, Sarbox has two potential effects: It increases the precision of the information provided by the CEO to investors, and it reduces aggregate corporate investment.

Our analysis has implications for corporate governance. Given the nature of internal organizational governance, boards are likely to choose CEOs from pools dominated by overconfident managers. While overconfidence benefits shareholders, excessive overconfidence does not, as it leads to overinvestment. An overconfident CEO also invests less in information acquisition, which further compromises investment decisions and the board's ability to judge them. This, however, is *not* a failure of corporate governance at the board level, but rather a consequence of constrained-efficient internal organizational governance. Moreover, Sarbox may reduce the incidence of CEO overconfidence.

Apart from the literature on corporate governance, our paper is also related to the economics of leadership (e.g., Hermalin (1998) and Van Den Steen (2005)), rank-order tournaments (e.g., Lazear and Rosen (1981) and Ramakrishnan and Thakor (1991)), overconfidence (e.g., Coval and Thakor (2005) and Gervais, Heaton, and Odean (2007)), the survival of overconfident agents (e.g., Kyle and Wang (1997), Van Den Steen (2004), and Wang (2001)), and managerial career concerns (e.g., Holmstrom (1999), Holmstrom and Ricart i Costa (1986), Milbourn, Shockley, and Thakor (2001), and Prendergast (1999)).

The rest of the paper is organized as follows. Section I presents the basic model. Section II analyzes the first- and second-best outcomes with and without promotion concerns. Section III examines how managerial overconfidence affects internal organizational governance. Section IV asks whether shareholders prefer an overconfident or a rational CEO. Section V examines corporate governance at the board level in terms of the board's decision to retain/fire the CEO, as well as the impact of Sarbox. Section VI summarizes the empirical predictions and Section VII concludes. All proofs are in the Appendix.

I. The Basic Model of Internal Organizational Governance

In this section, we describe the organization structure of the firm, the preferences of the agents involved, the probability distribution of project payoffs and their dependence on managerial ability, the compensation structure for agents, and the role of the CEO.

A. Organization Structure and Model Overview

We consider an all-equity firm with two levels of hierarchy. A CEO is at the top and $n > 2$ a priori identical managers report to her. The CEO determines the set of projects for which managers choose project risk levels. The CEO's choice is to either accept a new set of projects ("project portfolio" henceforth) or reject the new set in favor of existing operations with less uncertain aggregate payoff. The choice is made on the basis of the CEO's observation of a noisy signal—whose precision the CEO may be able to privately choose—of the payoff distribution of the project portfolio. The CEO's choice of the cut-off signal value above which she accepts the project portfolio is called the firm's "strategy."

Conditional on this strategy, each manager can choose the riskiness of his own project, and the aggregate payoff for the firm is the sum of the payoffs of individual managers' projects. The probability distribution for a particular manager's project payoff depends on four arguments: the CEO's ability, the manager's ability, the CEO's chosen strategy, and that manager's project risk choice.

We consider two periods. The most general version of the model can be described as follows. In the first period, the incumbent CEO, who has announced she will retire at the end of the first period, makes her strategy choice and individual managers choose their project risks. Individual managers' project payoffs, observed at the end of the first period, are used to infer abilities. The CEO chooses the highest-ability manager to succeed her and steps aside. This is the ability-filtering rule for promotions. The second period starts with the new CEO who selects and implements her strategy. Managers then make their risk choices and second-period project payoffs are observed.

In order to simplify the analysis, we suppress different features of the model when we analyze internal organizational governance and board governance. In Sections II and III, where we examine internal governance, we suppress the incumbent CEO's ability and first-period choices of strategy and signal precision. The payoff distribution for a manager's project thus depends only on that manager's ability and risk choice. In Section IV, we examine how a newly appointed CEO will behave in the second period, and thus we suppress the project risk choices and abilities of individual managers. The payoff distribution now depends only on the CEO's ability and strategy choice, and the CEO can choose the precision of the signal to implement her strategy. In Section V, we use the results of Section IV to examine board governance. We continue to suppress individual managers' abilities and risk choices, and also the CEO's signal precision choice. The question we ask is: If the incumbent CEO was not pre-destined to retire at the end of the first period, how would the board decide whether to retain or fire her after observing the first-period outcome? A key difference between Sections IV and V is that in Section IV the CEO knows she is playing an end game, whereas in Section V the board must recognize that the outcome observed at the end of the first period reflects a decision by a CEO who had career concerns.

B. Preferences

The shareholders are risk neutral. The CEO and the managers reporting to her are risk averse with von Neumann–Morgenstern utility u that is increasing and concave in compensation.

C. Abilities and Projects

The ability of manager i is A_i and that of the CEO is A_0. Managers do not know their own ability. The managers are ex ante observationally identical with independent and identically distributed abilities. The probability distributions of managerial ability and the CEO's ability are common knowledge. Each manager has one first-period project and chooses its risk. Manager i's project risk choice is $R_i \in [0, R_{\max}]$; this risk choice is unobservable to all but the manager himself. All projects require the same initial investment. The payoff y_i of a project managed by manager i is

$$y_i = \alpha + x_i, \tag{1}$$

where α, the "common component" of all projects' payoffs, depends on the set of projects and is stochastically independent of x_i, the "project-specific payoff." We assume α and x_i can be separately observed and define $X \equiv \{x_1, \ldots, x_n\}$. The probability density function of x_i, $f(x_i; R_i, A_i, A_0)$, depends on the project risk (which the manager chooses) and the abilities of the manager and the CEO (both unalterable attributes), with F being the corresponding cumulative distribution function. Since abilities are unknown, the probability density function of x_i unconditional on CEO and managerial ability, $f(x_i; R_i)$, is given by

$$f(x; R) = E[f(x; R, A_i, A_0)]. \tag{2}$$

The function f satisfies three requirements. First, it displays the monotone likelihood ratio property[4] with respect to managerial ability in that, for $A_L < A_H$, $f(x; R, A_L, A_0)/f(x; R, A_H, A_0)$ declines as x increases. Thus, a manager with higher ability has a higher probability of high payoffs. Second, the function $f(x_i; R_i)$ is "even"; $f(x; R) = f(-x; R)$.[5] Third, increasing project risk makes more extreme project payoffs more likely: If $R_H > R_L$, then $F(x; R_H) > F(x; R_L)$ for $x < 0$ and $F(x; R_H) < F(x; R_L)$ for $x > 0$.[6] The first and the third requirements are definitions of managerial ability and project risk. The requirement that $f(x_i; R_i)$ be "even" ensures that the effect of an increase in project risk is symmetric in making both high and low extreme payoffs more likely. The general specification of f ensures that our results do not depend on a particular distribution function. We give an example below that satisfies all three requirements (see Figure 1).

Example: The abilities A_i and A_0 are uniformly distributed over [–1/2, 1/2] and, given exogenous constants $k_0, k_1 > 0$,

$$f_i(y_i) = \begin{cases} \dfrac{1}{2R_i}\{1 + (y_i - \alpha)(k_0 A_0 + k_1 A_i)\} & \text{if} \quad |y_i - \alpha| \le R_i \\ 0 & \text{if} \quad |y_i - \alpha| > R_i. \end{cases} \tag{3}$$

D. Manager's Compensation

The exact form of the wage contract in the first period is unimportant for our results. It is natural to assume that the optimal wage contract for the first period will trade off the provision of project risk choice incentives against efficient risk sharing. Each manager also cares about his expected utility from future periods. This future expected utility will depend on his perceived ability at the end of the first period. A manager with a higher perceived ability has a higher expected project payoff and is therefore more valuable to the firm.

[4] See Milgrom (1981).

[5] In an earlier version of the paper, we had allowed mean project payoff to depend on project risk. All our results continue to hold with that specification. The current specification has project mean independent of risk to emphasize the results about risk-taking even in the absence of a risk-return tradeoff.

[6] If F_1 and F_2 are two distributions such that $F_2 - F_1$ has a single sign change from positive to negative, then F_2 is riskier than F_1 in the sense that if an agent with nondecreasing utility prefers F_1 over F_2, then so will a more risk-averse agent with nondecreasing utility (see Diamond and Stiglitz (1974) and Jewitt (1989)). This single-crossing property of distribution functions coupled with even distribution functions yields our characterization of risk. Higher risk in this sense also implies higher variance. And the converse is also true in special cases, for example, if distribution functions are restricted to the same family such as normal distributions with fixed means.

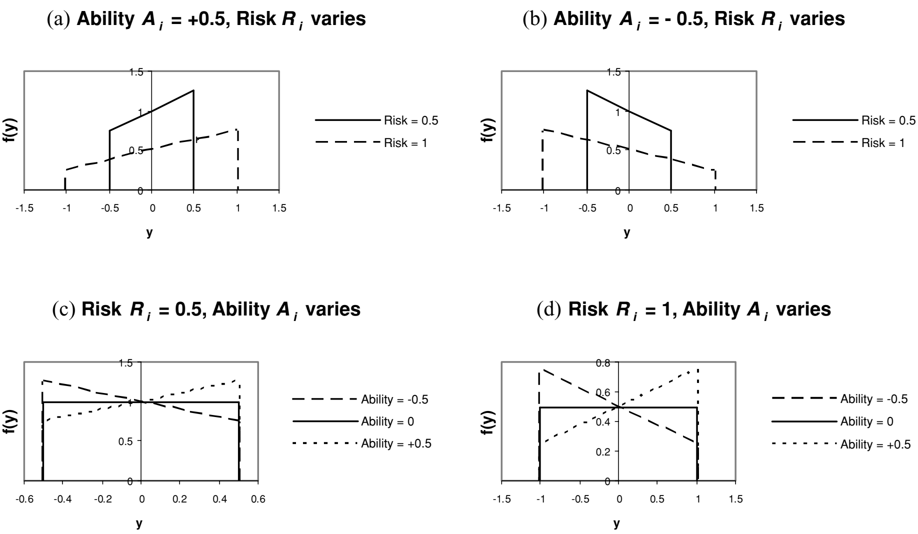

Figure 1. Probability density function of project payoff in the example ($\alpha = 0$, $A_0 = 0$, $k_1 = 1$).

Thus, in a competitive labor market, managerial wages will be increasing in perceived abilities. While all managers are initially observationally identical, the first-period payoffs induce a revision of beliefs about managerial ability. Project-specific payoffs X observed at the end of the first period permit distinct ability reassessments.

From the optimal dynamic contracting literature (e.g., Rogerson (1985)), we know that the dependence of compensation on perceived ability will be strongest if managers are risk neutral, and weaker but nonetheless positive, with risk-averse managers.[7] An optimal contract can be derived by taking into consideration the risk aversion of managers, the value of perceived ability in the labor market in current and future periods (which itself derives from the dependence of future project payoffs on ability), and the time value of money. We do not solve for the optimal compensation contract here, but assume that manager i's compensation (for current and future periods) is an increasing function $w(E[A_i])$ of the perceived ability of the manager at the end of the first period.

E. The Role of the CEO

The CEO leads by determining corporate strategy and thereby setting an overall direction for the firm. Thus, while managers make decisions that

[7] This effectively means that managerial compensation will be increasing in a manager's first-period project-specific payoff. The compensation may also depend on the project-specific payoffs of other managers, but will be independent of the common component of the project payoff as the manager has no control over that.

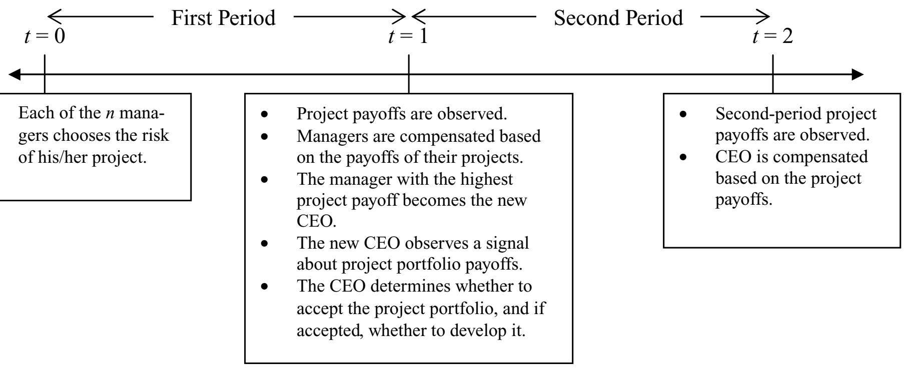

Figure 2. Sequence of events.

affect the prospects of their individual projects, the CEO's decision affects the prospects of *all* the projects, that is, the firm as a whole.

We model this in a simple way. The common component α of the project payoff y_i in (1) depends on the set of projects determined by the CEO's choice of strategy S. Further, the CEO's ability A_0 affects the probability density function f of the project-specific payoff x_i. In subsequent sections, we will view S as the CEO's choice of accepting or rejecting a project portfolio. This decision will be based on a noisy but informative signal of portfolio quality, whose precision the CEO may be able to control. We will introduce these elements in greater detail in Sections IV and V. For now, it suffices to note that $\partial F(y_i \mid .,.,.)/\partial A_0 < 0 \; \forall \, i$, where F is the cumulative distribution function associated with density function f. Thus, higher ability results in a stochastically higher payoff for each project. The aggregate firm payoff, $Y = \sum_{i=1}^{n} y_i = n\alpha + \sum_{i=1}^{n} x_i$, consists of two independent components: the "base payoff" $n\alpha$, which is the sum of n common components of project payoffs, and $\sum_{i=1}^{n} x_i$, the sum of all project-specific payoffs.

This specification will allow us to verify later that it is subgame perfect to promote the manager with the highest perceived ability to CEO in the second period. We assume that there is a private benefit $B > 0$ to the manager from being appointed CEO. The sequence of events is summarized in Figure 2.

II. The Impact of Internal Organizational Governance on Managerial Project Risk Choices: The First Best and Promotion Concerns

In this section, we suppress the second-period game in which a new CEO is at the helm, and focus on first-period project risk choices. Each manager's risk choice depends on his compensation contract and on the distribution of his project-specific payoff x_i, but not on the unknown abilities of the CEO and managers or the CEO's strategy S. Thus, we return to the functional form in (2) for this analysis. In this section, all managers are assumed to be rational. We begin by examining the risk choices of managers without promotion concerns and then introduce promotion concerns.

A. First-Best Risk Choices

Suppose each manager's ability is known. Managerial compensation will still depend on perceived ability, but will be independent of the project payoff because there is no moral hazard. Each manager will act in the shareholders' interest by choosing project risk to maximize expected project payoff. If project risk and managerial ability interact in determining the expected payoff, managers with different abilities may make different risk choices. In the Example in the previous section, the expected payoff of a project managed by manager i can be computed to be $k_1 A_i R_i^2/3 + \alpha$. To maximize this, a manager with ability $A_i > 0$ chooses the maximum risk $R_{\max}$ and a manager with lower ability chooses minimum risk of zero. The risk choice is increasing in the manager's ability because the benefit of managerial ability increases with project risk. In an alternative specification where managerial ability and project risk are separable, all managers will choose the same risk regardless of their abilities.

B. Second-Best Risk Choices without Promotion Concerns

Suppose each manager's unknown ability is inferred from his project payoff. The manager's project risk choice is also unobservable. For now, we suppress the implicit tournament that determines the new CEO at the end of the first period. Thus, each manager chooses his project risk seeking only to maximize the expected utility of compensation, taking as given his wage contract, $w(y_i)$, in which compensation is an increasing function of the project payoff. Of course, the optimal contract in equilibrium is based on rational anticipation of managerial risk choices in response to the wage contract.

Each ex ante identical manager's expected utility calculation is unconditional on unknown ability and is identical across managers. The expected utility of manager i with risk choice R_i is

$$U_i = E[u(w_i(X)) \mid R_i, \hat{R}], \tag{4}$$

where $\hat{R}$ is manager i's belief about the risk choices of the other managers. Since all managers solve the same problem, they choose the same risk R^o.

The choice of risk by the manager depends on the manager's risk aversion as well as the form of the wage function. If the wage is a linear function of posterior managerial ability, increasing project risk increases the risk of the wage without altering the mean wage. This reduces the expected utility of the risk-averse manager. If the wage function is concave (convex) in posterior managerial ability, increasing project risk decreases (increases) the mean wage. These two considerations determine the optimal risk choices of managers in the absence of career concerns. Without explicitly solving for this risk choice, we shall focus on how promotion concerns affect this risk choice.

C. Internal Organizational Governance and Promotion Concerns

We now assume that the managers compete for promotion to CEO at the end of period one, so our focus is on the first-period project choices. The current

CEO wishes to promote the person with the highest ability. Since all managers are ex ante identical and the probability density of the project payoff follows the monotone likelihood ratio property, the manager with the highest project payoff has the highest perceived ability after the first period and is promoted to CEO for the second period.

LEMMA 1: *If manager i chooses risk greater than the risk chosen by all the other managers in the first period, his probability of promotion is $P_i > 1/n$, where $1/n$ is the probability of promotion for any manager if all the managers choose the same risk.*

Lemma 1 says that a manager can improve his promotion prospects by increasing his project risk. The intuition is as follows. If the manager chooses the same risk as the other managers, he has the same project payoff distribution at the end of the first period. Thus, with n managers who are a priori identical, the probability of being promoted is $1/n$. The choice of a riskier project by manager i increases his promotion probability because when manager i does have a high payoff, his payoff tends to be higher than that of all the competing managers. The reason is that higher risk makes more extreme payoffs more likely even when it does not affect the mean. Because risk choices are unobservable, the effect of risk on the project payoff cannot be disentangled from that of ability. Thus, a payoff higher than those of competing managers makes manager i appear more able than his competitors and gets him promoted.[8]

To illustrate this, suppose we have 20 managers and four payoff states: very low (VL), low (L), high (H), and very high (VH). With moderate risk, the payoff probabilities are 0.49 each for L and H and 0.01 each for VL and VH. With high risk, each payoff state has equal probability. Now, if all managers choose moderate risk, the probability of being promoted to CEO is 1/20 for each. A deviating manager who chooses high risk has a higher probability of achieving VH and thus a higher chance of being promoted to CEO, even if the expected project payoff is higher with moderate risk.

The tournament for promotion to CEO causes each manager to base his project risk choice on the expected risk choices of other managers. To ensure the existence of a pure-strategy equilibrium, we assume that a manager's best-response project risk choice is a continuous function of the project risk choices of other managers, bounded below the maximum risk $R_{\max}$. We now have the following result.

[8] This intuition does not work when there are only two managers because then manager i's project payoff is compared with the project payoff of the other manager, which is equally likely to be above or below the mean. By increasing risk, manager i makes moderate payoffs less likely and very high or very low payoffs more likely. This increases his probability of promotion conditional on the other manager achieving a high payoff but decreases his probability of promotion conditional on the other manager achieving a low payoff. When there are more than two managers, manager i's payoff is compared with the maximum of the project payoffs of the remaining managers and this maximum is more likely to be above the mean payoff than below the mean payoff. By increasing the likelihood of a very high payoff, manager i increases his chance of exceeding the maximum payoff of the remaining managers.

Proposition 1: *When managers are competing to be CEO, they choose projects with risk R^* that is greater than R^o, the risk chosen when the managers have no promotion concerns and seek only to maximize expected utility of compensation, and R^* is increasing in the benefit from promotion (B).*

This proposition follows readily from Lemma 1. Since the manager can enhance his promotion probability by choosing more risk, in a Nash equilibrium, all managers choose higher risk than they would in the absence of promotion concerns. Risk thus has both a cost and a benefit for the manager. The personal cost of the reduced expected utility from riskier compensation is traded off against the benefit of an enhanced promotion probability. As long as $B > 0$, promotion concerns increase the chosen risk. Since all managers are identical, they all choose projects with identical higher risk.[9]

III. The Effect of Managerial Overconfidence on Internal Organizational Governance

We now add overconfidence to the analysis and ask: Are overconfident managers more or less likely than rational managers to succeed in becoming CEO? We begin by examining first-period risk choices and the relative promotion probabilities of overconfident and rational managers. We find that an overconfident manager has a higher probability of being promoted when nobody realizes there is an overconfident manager in their midst. We then ask what happens if every manager recognizes that there is some probability that he is overconfident. The previous result is qualitatively sustained.

Suppose one of the managers, say manager i, is overconfident. By this we mean that the manager assigns too narrow a confidence interval when forecasting the project payoff, thereby underestimating the risk of the project payoff. Much of the evidence on overconfidence corresponds to agents overestimating the precision of their forecasts, a phenomenon that is also termed "hyperprecision" (see Kahneman, Slovic, and Tversky (1982), p. 475). Another notion of overconfidence used in the literature corresponds to agents overestimating the precision of an exogenous noisy signal (e.g., Daniel, Hirshleifer, and Subrahmanyam (1998)). Our definition of overconfidence is related to both these notions of overconfidence if the manager's estimate of project payoff is based on a signal, so that by overestimating the precision of the signal, the manager also overestimates the precision of the forecast of the project payoff.[10]

[9] The notion of excessive risk-taking has a parallel in the R&D literature. Bhattacharya and Mookherjee (1986), Dasgupta and Maskin (1987), and Klette and de Meza (1986) consider the choice of research strategies by competing firms where research strategies affect the probability distribution of invention times or quality of research output. These papers show that with a "winner-take-all" payoff structure, the research strategies chosen in a market equilibrium are riskier than those in a socially optimal outcome.

[10] Two other notions of overconfidence used in the literature are overestimating one's ability to accomplish something and overestimating the probability of good outcomes (optimism). The results of Section IV show that a CEO who is overconfident in the sense that she overestimates the precision of her information will be too optimistic about the project portfolio she accepts when she bases the portfolio acceptance decision on her information.

If a project has risk R, the overconfident manager erroneously believes that the project's risk is R/C, where $C > 1$ is the degree of overconfidence. We assume that no one (including manager i) suspects that manager i is overconfident. Thus, all managers aim to choose risk R^* in the first period (Proposition 1). The overconfident manager, believing he is choosing risk R^*, inadvertently chooses CR^*. This leads to the following proposition.

PROPOSITION 2: *An overconfident manager is more likely to get promoted at the end of the first period when no one realizes the problem of overconfidence. The relative increase in the probability of the overconfident manager's promotion is increasing in the degree of his overconfidence.*

We saw in Proposition 1 that the risk aversion of managers and the prospects of promotion interact to determine the optimal level of risk chosen by each manager. An overconfident manager unwittingly chooses a project with higher risk and ends up hurting himself in an expected utility sense. The higher risk, however, increases the probability of promotion for the overconfident manager. The strength of this effect is increasing in the degree of overconfidence.

We now assume that all managers are ex ante identical and each is likely to be overconfident with probability π. The occurrence of overconfidence is independent across managers. The degree of overconfidence is C for an overconfident manager. This information is common knowledge, which means that each manager recognizes that the probability with which he is overconfident is π.

All managers face the same problem of maximizing their expected utility over compensation and promotion. This problem is more complicated now because each manager must consider the possibility that he may be rational or overconfident. At the same time, he must realize that others may also be rational or overconfident in a similar way. Solving this problem, we get the following result:

PROPOSITION 3: *When each manager is likely to be overconfident with probability π and this is common knowledge, all managers believe they have chosen risk R^{**} in the first period. For a rational manager, the truly chosen project risk is also R^{**}, while an overconfident manager's project has risk CR^{**}. An overconfident manager is more likely to get promoted than a rational manager.*

Because managers are ex ante identical, it is not surprising that all the managers adopt the same strategy and choose those projects that they think have risk R^{**}. They realize that the true risk of their project could be R^{**} or CR^{**}, depending on whether they are rational or overconfident. Thus, introducing overconfident managers results in cross-sectional variation in the risks of projects.

Since overconfident managers choose riskier projects, they are more likely to get promoted than rational managers. Thus, the posterior probability that the promoted manager is overconfident is higher than π, the ex ante probability of overconfident managers. Surprisingly, this does *not* imply that the promotion rule is inefficient for shareholders. The ability-filtering rule used for promotions is rational in that the manager with the highest project payoff is still the most likely to have the highest ability. The fact that overconfident managers are more likely to get promoted does not mean that average ability will be lower for promoted managers than for those not promoted; the reason is that ability and overconfidence are uncorrelated.[11] Since this effect arises from a promotion rule that rewards the manager with the highest perceived ability, our analysis implies that a greater emphasis on merit-based promotions makes it more likely that overconfident managers are appointed CEOs.

Thus far, we have derived our results using a two-period model. Similar results will obtain, however, in *any* finite-horizon model as long as project choices are made in only one period. That is, we can imagine a model in which there is a CEO at the helm with n managers reporting to her. The CEO will be in office for $T \in (1, \infty)$ periods.[12] Managers make project risk choices at the start of the first period. Each manager's choice yields identical and independently distributed (i.i.d.) payoffs over each of the T periods that the current CEO is in place; project risk choice is unalterable for T periods once it is made in the first period. After T payoffs have been observed on each project, there is a posterior assessment of managerial ability, based on which a new CEO is chosen. The new CEO makes a one-time strategy choice, which is followed by a second (and final) project risk choice by each manager. Beyond this point, project payoffs can be observed for any finite number of time periods. Such a structure will yield essentially the same results we have at present. What seems intractable is a model in which the CEO makes multiple strategy choices and managers make multiple project risk choices.

IV. The Efficiency of Internal Organizational Governance and Its Impact on Board Governance: Is an Overconfident CEO Good for the Firm?

Given that internal organizational governance favors an overconfident manager as CEO, we now ask whether this benefits shareholders more than having a rational CEO. That is, we focus on the second-period behavior of a newly appointed CEO who screens a project portfolio for investment on the basis of a private signal whose precision the CEO may be able to choose.

In Subsection A, we show that a rational risk-averse CEO underinvests in new project portfolios. In Subsection B, we show that an overconfident CEO is better for shareholders than a rational CEO as long as she is not too overconfident; excessive overconfidence diminishes firm value. Subsection C shows that overconfidence differs in effect from low risk aversion. Lower CEO risk aversion *always* benefits the shareholders, whereas extreme CEO overconfidence

[11] It would be interesting to extend the analysis to see how the choice of riskiness and the promotion probability are affected if managerial ability and the probability of being overconfident are correlated. Even when ability and overconfidence are independent across managers, the promotion process described here would lead to a negative correlation between ability and overconfidence for promoted CEOs. The reason is that the promoted CEO's high project payoff can be attributed either to high ability or to high risk-taking due to overconfidence.

[12] We ignore the possibility of firing the CEO during these periods.

can hurt the shareholders. Subsection D establishes another important difference, namely, that the Revelation Principle can be used to sort out managers with different degrees of risk aversion but *not* overconfidence. Finally, in Subsection E, we show that an overconfident CEO underinvests in information precision, and yet is preferred if moderately overconfident.

A. A Model of Underinvestment

Recall that the aggregate payoff Y consists of the base payoff $n\alpha$ and the project-specific payoffs X. The CEO's strategy S affects the base payoff while the CEO's ability affects the project-specific payoffs for all projects. Let the CEO's strategy choice be $S \in \{accept, reject\}$, that is, the CEO can either accept or reject a new project portfolio. If $S = accept$, then individual managers invest in projects from the new portfolio. If $S = reject$, the managers invest in projects from existing operations. In each case, the managers have payoff risk choices to make at their level. The CEO's decision is made on the basis of a noisy but informative *private signal* s of "project portfolio quality" (defined below). We take the precision of s as fixed for now, but in Subsection E we will let the CEO choose this precision, and, we assume that there is a board acting in the shareholders' interest that designs the CEO's compensation.

The project portfolio differs from the existing operations in that it requires development and has greater aggregate payoff risk. The CEO can make an initial investment in project portfolio development through research and information acquisition as well as by spending time on implementation. Project portfolio development imposes a disutility of c on the CEO. A project portfolio developed by the CEO turns out to be good with probability p and bad with probability $1 - p$, where p, referred to as the *project portfolio quality*, is a random variable. The base payoff $n\alpha$ equals l for a bad project portfolio and h for a good project portfolio, where h and l are constants with $h > l$. If the CEO accepts the project portfolio but does not develop it, the portfolio will be bad with probability one. This can also be interpreted as the CEO rejecting the portfolio and investing instead in value-depleting activities that the board cannot distinguish from the portfolio. Finally, the CEO can reject the project portfolio, in which case the base payoff is r, where $r \in (l, h)$ is a constant. We let ω indicate the "project portfolio outcome": $\omega = H$ if the portfolio is accepted and turns out to be good (base payoff h), $\omega = L$ if the portfolio is accepted and turns out to be bad (base payoff l), and $\omega = R$ if the portfolio is rejected (base payoff r).

A manager's project-specific payoff depends on the manager's ability, the CEO's ability, and the manager's risk choice. The manager's risk choice will be independent of the project portfolio outcome ω because managerial compensation is independent of ω as noted earlier. Thus, project-specific payoffs are also independent of ω. Since managers are ex ante identical, the project-specific payoffs are independently and identically distributed. The probability distribution of the project-specific payoff x_i, unconditional on the manager's ability and risk choice, $\xi(x_i, A_0)$, depends on the CEO's ability A_0. The CEO's ability is unknown and only its probability density function, ψ, is known, so the distribution $\xi(x_i)$

is obtained by taking the expectation of $\xi(x_i, A_0)$ over the CEO's ability, A_0. The probability distribution ξ follows the monotone likelihood ratio property with respect to the CEO's ability; $\xi(x_i, A_H)/\xi(x_i, A_L)$ is increasing in x_i for $A_H > A_L$. Thus, the CEO's ability affects the aggregate payoff in all states, but *not* the project portfolio quality (or "success" probability).

There are two forms of information asymmetry between the CEO and the board. First, the board cannot directly verify whether the CEO developed the portfolio that was accepted. It can noisily infer the CEO's action by observing the portfolio outcome ω. An $\omega = H$ outcome reveals that the CEO developed the portfolio, but $\omega = L$ may result even if the CEO developed the portfolio. Second, the CEO privately observes the portfolio quality signal s before deciding whether to accept the portfolio. The board does not observe s, but its probability density function is common knowledge.

We now specify how the CEO updates her beliefs about the portfolio quality p based on the signal s. The project portfolio quality p is uniformly distributed over [0, 1]. With probability q^*, the signal s equals p, and with probability 1 – q^* it is uninformative about p and uniformly distributed over [0, 1].[13] We can view q^* as the precision of s. The expected value of p, conditional on observing the signal s, is given by

$$E\ [p \mid s] = sq^* + 0.5(1 - q^*). \tag{5}$$

The posterior estimate of p is thus a weighted average of the observed signal s and its prior mean 0.5. The board, however, knows that the CEO may be overconfident and that this may bias the CEO's beliefs about the precision of the signal s. A CEO with confidence C believes that s equals p with probability $q(C)$ and is uninformative with probability $1 - q(C)$, where q is an increasing function and $q(1) = q^*$. This means that an overconfident CEO's ($C > 1$) estimate of the precision of s exceeds the true precision q^*, and the project portfolio quality is incorrectly estimated to be

$$p^C = p(s, C) = sq(C) + 0.5[1 - q(C)] \tag{6}$$

while the rational estimate is $p = p(s, 1)$. This manifestation of overconfidence is similar to that in Section III, where an overconfident manager overestimates the precision of his information about the project payoff and consequently underestimates payoff risk; here, an overconfident CEO overestimates the precision of her private signal about project portfolio quality and thus underestimates the risk of the signal. In both instances, overconfidence is synonymous with overestimating the precision of one's information. The degree of overconfidence of the CEO, C, is unknown and only its distribution $\mu(C)$ is common knowledge. Thus, we do not preclude a diffident CEO with $C < 1$.

The CEO does not believe that she is overconfident. Rather than viewing C as a measure of overconfidence, she interprets $q(C)$ as the precision of the signal s; this precision depends on the parameter C that she will learn after entering

[13] All results follow if the common range of p and s is restricted to a proper subset of [0, 1].

into the wage contract but before making the investment decision. The CEO shares common beliefs about the probability distribution $\mu(C)$, so, despite her overconfidence, she correctly anticipates how she may interpret the signal in the future.[14]

The CEO's wage is W_R if $\omega = R$, W_L if $\omega = L$, and W_H if $\omega = H$. It is independent of the project-specific payoffs x_i because the CEO's strategy does not affect x_i. The CEO's opportunity wage is W_0, so she needs a minimum expected utility of $u(W_0)$ to participate. The board designs the contract to incent the CEO appropriately for accepting or rejecting a portfolio and for developing accepted portfolios.

Suppose the board wants the CEO to reject the portfolio when assessed portfolio quality $p(s, C) < p^*$ and accept and develop the portfolio when $p(s, C) \geq p^*$.[15] Then the board's problem is

$$\underset{p^*, W_H, W_L, W_R}{\text{Max}} \left\{ \begin{array}{l} E\left[\sum_{i=1}^{n} x_i\right] \\ + \Pr(p(s, C) \geq p^*) \times E\left[p(s, 1)(h - W_H) + \{1 - p(s, 1)\}(l - W_L) \mid p(s, C) \geq p^*\right] \\ + \Pr(p(s, C) < p^*) \times (r - W_R) \end{array} \right\} \tag{7}$$

subject to

$$p^* u(W_H) + (1 - p^*) u(W_L) - c = u(W_R) \tag{8}$$

$$W_H \geq W_L, \tag{9}$$

$$W_R \geq W_L, \tag{10}$$

and

$$\begin{aligned} &\Pr(p(s, C) \geq p^*) \times E[p(s, C) u(W_H) + \{1 - p(s, C)\} u(W_L) - c \mid p(s, C) \geq p^*] \\ &\quad + \Pr(p(s, C) < p^*) u(W_R) \geq u(W_0). \end{aligned} \tag{11}$$

The objective in (7) is the expected payoff to shareholders net of the CEO's compensation. There are three terms. The first is the expected value of the sum of project-specific payoffs. The second term is the probability of portfolio acceptance multiplied by the expected value of the base payoff net of the CEO's wage conditional on portfolio acceptance: $h - W_H$ when $\omega = H$ (an event with probability $p(s, 1)$), and $l - W_L$ when $\omega = L$ (an event with probability $1 - p(s, 1)$). The third term is the probability of portfolio rejection multiplied by the payoff $r - W_R$ when $\omega = R$. Constraints (8) to (10) are incentive compatibility

[14] This means she believes she will act optimally and does not realize that the investment rule based on her biased beliefs will not maximize her expected utility.

[15] A contract that induces the CEO to accept a portfolio without developing it will induce the CEO to accept all portfolios and will not be optimal as long as a portfolio can have sufficiently low NPV (p can be sufficiently low).

constraints, where (8) and (9) ensure that the CEO prefers accepting and developing a portfolio to rejecting it if her estimate of portfolio quality is at least p^*, and prefers to reject otherwise. Constraint (10) ensures that the CEO prefers rejecting a portfolio to accepting it but not developing it. The CEO's participation constraint is given by (11). The first term in the CEO's expected utility is the expected utility from her wage net of the development cost conditional on portfolio acceptance and the second term is the expected utility from her wage conditional on portfolio rejection.[16]

PROPOSITION 4: *If the deviation in the CEO's confidence from rationality is sufficiently small, the second-best equilibrium induces underinvestment, so that when the CEO is indifferent between accepting and rejecting a project portfolio, the shareholders are strictly better off if the portfolio is accepted.*

The intuition is that the board, unable to observe whether the CEO develops a portfolio, offers the CEO payoff-contingent compensation to incent her to develop an accepted portfolio at a personal cost. However, the resulting compensation risk causes the risk-averse CEO to prefer the certain wage from portfolio rejection to an equal *expected* wage from portfolio acceptance, causing underinvestment.

The optimal compensation contract trades off the cost of underinvestment and a high expected wage when the compensation is more payoff-contingent against the cost of the CEO not developing the portfolio if the compensation is less payoff-contingent. The resulting equilibrium allows some underinvestment by the CEO. That is, the shareholders' expected payoff net of the CEO's compensation is increased if the CEO lowers the minimum portfolio quality for acceptance below p^*.

We now seek to determine the set of project portfolios whose acceptance benefits shareholders. It is straightforward to see that the shareholders' expected incremental payoff from accepting a portfolio is monotonic in the assessed portfolio quality p. Whether it is increasing or decreasing in p depends on the managerial wage contract $\{W_H, W_L\}$. It is increasing in p if the wage contract is such that $h - W_H > l - W_L$, so that the shareholders' expected payoff is higher with a good portfolio than with a bad portfolio. It is decreasing in p if $h - W_H < l - W_L$. However, optimal risk-sharing between the shareholders and the CEO rules out wage contracts of the latter kind because they necessarily imply $W_L > W_H$ (since $h > l$), and thus a negative relationship between the shareholders' and the CEO's payoffs. We omit the proof because it resembles that of Proposition 4 in Grossman and Hart (1983). The shareholders' expected incremental payoff from accepting a portfolio is thus increasing in p, and it is easy to verify that the portfolio quality p^{**} at which the shareholders are indifferent between acceptance and rejection is

[16] The board's expected payoff assessment is completely rational. It values projects based on rational beliefs about project quality, and it accounts for the fact that the CEO's investment decision may be based on a biased project-quality inference. However, the expected utility of the CEO in the individual rationality constraint computes the values of accepted projects that will be accepted based on the CEO's biased project quality inference.

$$p^{**} = \frac{(r - W_R) - (l - W_L)}{[h - W_H] - (l - W_L)}. \tag{12}$$

We can now characterize the set of portfolios that shareholders would like the CEO to accept.

COROLLARY 1: *There is a probability value $p^{**} < p^*$ such that any portfolio with $p^{**} < p < p^*$ is rejected by the CEO, but the shareholders' expected payoff is greater with portfolio acceptance than with rejection.*

Next, we examine the impact of the CEO's ability on firm value.

LEMMA 2: *Firm value is increasing in the CEO's ability, A_0.*

The intuition is that a higher-ability CEO is more likely to generate higher project-specific payoffs x_i. This result verifies that it is subgame perfect to appoint as second-period CEO the manager with the highest perceived ability.

B. Effect of CEO Overconfidence on Investment

We now consider how CEO overconfidence affects investment policy and firm value. We assume that the exogenous parameters are such that the shareholders will not want a rational CEO to accept a project portfolio without observing the signal. Specifically, we assume that $p^* > p(0.5, 1)$. If we define signal values s^* and s^{**} to correspond to the probabilities p^* and p^{**}, then

$$s^* > 0.5. \tag{13}$$

PROPOSITION 5: *The value of the signal above which an overconfident CEO accepts a project portfolio decreases as her degree of overconfidence (C) increases. There is a threshold degree of overconfidence C^* such that firm value increases with the degree of CEO overconfidence for $1 < C < C^*$, but decreases with the degree of CEO overconfidence for $C > C^*$.*

The intuition is as follows. The CEO accepts portfolios only when she observes a signal high enough to ensure that the benefit of a high probability of the high payoff outweighs the cost of bearing risk in accepting the portfolio. There is a cutoff value of the signal, s^*, below which the rational CEO does not accept a portfolio. However, when an overconfident CEO observes a signal just below s^*, she overestimates the "good news" in the signal—the probability of the high payoff—and consequently accepts the portfolio. The greater the CEO's overconfidence, the lower is the cutoff signal at which the CEO accepts the portfolio. The set of portfolios that are accepted grows with CEO overconfidence.[17] This is

[17] This irrational behavior reduces the CEO's expected utility because she is accepting project portfolios with risks that are too high relative to her compensation for bearing them.

consistent with Malmendier and Tate (2008), who document that overconfident CEOs overinvest in acquisitions and the market reacts more negatively to such acquisitions.

Interestingly, the impact of CEO overconfidence on firm value is *nonmonotonic*. Proposition 4 shows that a rational, risk-averse CEO underinvests. An overconfident CEO accepts all portfolios a rational CEO accepts and also invests in additional value-enhancing portfolios that a rational CEO rejects. Thus, CEO overconfidence attenuates some of the underinvestment inefficiency due to CEO risk aversion. This is the driving force behind Proposition 5. As the CEO's overconfidence increases, she is also willing to invest in portfolios that have lower probabilities of the high payoff and produce smaller marginal increases in shareholders' wealth. Thus, firm value increases with CEO overconfidence at a decreasing rate up to a point, and eventually it starts declining as overconfidence causes the CEO to accept portfolios that even the shareholders want rejected. That is, while moderate overconfidence diminishes underinvestment inefficiency, higher levels of overconfidence create overinvestment.

The result in Proposition 5 that moderate CEO overconfidence enhances firm value depends on the assumption that the CEO is risk averse. If the CEO is risk neutral or risk loving, any level of overconfidence will lead to excessive risk and reduce firm value.

Proposition 5 also suggests that the value of CEO overconfidence to the shareholders will be circumstance-specific. Since the underinvestment inefficiency due to CEO risk aversion will be greater in riskier industries, CEO overconfidence will be more highly valued in such industries. Put a little differently, the threshold level of confidence above which CEO overconfidence will adversely affect firm value will be higher for firms in riskier industries, controlling for CEO risk aversion.

Gervais, Heaton, and Odean (2007) find results that complement Proposition 5. They show that overconfidence and optimism can be an alternative to executive stock options in resolving agency problems, and thereby enhancing firm value. They emphasize that executive compensation design should take into account executive behavioral biases, and that executive stock options that are efficient with managerial rationality may induce excessive risk taking with managerial overconfidence.

C. Differences in Effects of Reduced CEO Risk Aversion and Overconfidence on Investment

Because the effect of overconfidence on reducing underinvestment is similar to that of a reduction in the risk aversion of the CEO, it may appear that overconfidence and reduced risk aversion are isomorphic. We now show, however, that this intuition is not correct.

PROPOSITION 6: *Firm value is monotonically decreasing in the risk aversion of a rational CEO.*

The intuition is as follows. As shown in Proposition 4, with payoff-contingent compensation, the risk-averse CEO's preference becomes skewed towards rejecting project portfolios. This reduces firm value. Proposition 6 shows that as risk aversion decreases, the cost of bearing risk also decreases and it is cheaper for the shareholders to incent the CEO to accept good portfolios. This increases firm value.

Propositions 5 and 6 highlight one distinction between the effects of low risk aversion and overconfidence. While both help ameliorate underinvestment inefficiencies, overconfidence may also induce overinvestment because of the bias it introduces in beliefs, so firm value is *not* monotonically increasing in CEO overconfidence. A decrease in risk aversion, by contrast, is unambiguously beneficial because it enables contracts with greater payoff-dependence and better incentives.

Like Proposition 5, Proposition 6 relies on the CEO being risk averse. There is considerable empirical evidence that managers are risk averse. Amihud and Lev (1981) hypothesize that conglomerate mergers are undertaken by risk-averse managers to reduce their undiversified employment risk, and provide supporting empirical evidence. Grund and Sliwka (2006) use the German Socioeconomic Panel, which contains data on individual risk aversion, to present evidence on the link between risk aversion and pay-for-performance sensitivity in executive compensation. Khambata and Liu (2006) use data for 14 countries to show that firms in countries with higher risk aversion exhibit lower dividend ratios and lower propensity to pay dividends. Pattillo and Söderbom (2000) find evidence of the impact of managerial risk aversion on the decisions of Ghanaian manufacturing firms.[18] Finally, Suto, Menkhoff, and Beckmann (2005) uncover evidence of risk aversion among Japanese and German fund managers.

There is also ample evidence of overconfidence. Larwood and Whittaker (1977) study a sample of corporate presidents and find them to be unrealistic in their predictions of success (consistent with overestimating precisions of private signals). Weinstein (1980) provides evidence that people are especially overconfident about the projects to which they are highly committed, which supports our assumption that some managers are overconfident about their own projects. More recently, Ben-David, Graham, and Harvey (2007) present empirical evidence that managerial overconfidence is associated with aggressive corporate policies, including investments, financing, and executive compensation. Thus, it appears that overconfidence provides a much more plausible foundation for preferences that can lead to value-destroying overinvestment than do risk-seeking preferences.[19]

An interesting empirical issue is whether risk aversion and overconfidence are correlated in the cross-section, so that the more overconfident CEOs are also less risk averse. In this case, the combination of increased overconfidence and diminished risk aversion will positively impact firm value, generating a

[18] Risk aversion in the studies cited here impacts corporate decisions that typically involve the CEO and other high level executives. That is, this is *not* merely evidence of risk aversion in the general population, but rather risk aversion among CEOs as well as managers who represent the set of agents from whom CEOs are selected.

[19] Clearly, the value of CEO overconfidence to the shareholders depends on firms being unable to fully solve the underinvestment problem by selecting appropriately *risk-tolerant* CEOs.

result qualitatively similar to ours, but as one approaches risk neutrality, overconfidence will be unambiguously bad. That is, the cutoff point above which overconfidence will decrease firm value will be lower in this case than in our model. Moreover, with a sufficiently high cross-sectional correlation, most of the presumed effect of overconfidence would be captured by risk tolerance in empirical tests, and overconfidence would have little *independent* effect. While we are not aware of any direct test of the hypothesis that overconfidence is positively cross-sectionally correlated with risk tolerance, there is illuminating evidence that shows that risk aversion is pervasive and yet overconfidence exerts an independent effect on outcomes that cannot be explained by risk aversion alone. For example, Beetsma and Schotman (2001) use data from a Dutch television show, Lingo, to present evidence of risk aversion, but they also find that the predictions fit the data much better when overconfidence is added to the model. Guiso and Jappelli (2006) use data on the customers of a leading Italian bank to show strong evidence of overconfidence among individual investors, but more importantly, they also show that their findings *cannot* be explained by a possible correlation between risk aversion and overconfidence. Grinblatt and Keloharju (2006) use data on Finnish investors and find that overconfidence has a large effect on the investor's decision of whether to trade and that the results are not driven by differences in investors' risk aversion. Finally, Barber and Odean (2001) document that men exhibit greater overconfidence in trading behavior than women do, and that differences in risk aversion do not explain this.

D. Another Difference between Overconfidence and Reduced Risk Aversion: Sorting Agents with Unobservable Cross-sectional Heterogeneity in Overconfidence and Risk Aversion

Another way to appreciate the difference between overconfidence and reduced risk aversion is to recognize that managers are likely to know their preferences but not whether they are overconfident, since recognition of their own overconfidence would cause these managers to become rational; remember that overconfidence leads to a lower expected utility. So, if the shareholders can identify managers by their degrees of risk aversion, they will use this information in the optimal promotion rule for two reasons. First, this information helps shareholders separate the effect of a manager's risk choice on his project payoff from the effect of his ability. Second, managerial risk aversion may be another attribute in addition to perceived ability in the choice of the CEO. Moreover, second-period contracts will be customized to the known risk aversion of the new CEO, unlike the CEO's unknown overconfidence.

To show that shareholders can get the managers to reveal their risk aversion but not their overconfidence, we extend the first-period model to allow for cross-sectional variation in overconfidence as well as risk aversion. Suppose manager i has von Neumann–Morgenstern utility u_{σ_i} over his compensation, w, where σ_i measures the risk aversion of the manager. Higher σ means greater risk aversion in the sense that

$$\frac{d}{d\sigma}\left(-\frac{u''_{\sigma}(w)}{u'_{\sigma}(w)}\right) > 0. \tag{14}$$

The risk aversion σ_i is privately known to manager i while shareholders only know that $\sigma_i \in \Lambda$. Of the many potential risk aversion revelation mechanisms available to shareholders, we shall examine the one in which each manager reports his risk aversion and receives a compensation contract from a menu of contracts with varying degrees of risk sharing. Our objective here is to uncover an implementable revelation mechanism rather than solve for the most efficient mechanism. From the Revelation Principle, we can restrict our attention to direct revelation mechanisms in which the manager reports her risk aversion *truthfully* and accepts the pre-announced wage contract associated with that report.

Each manager reports his risk aversion and privately chooses project risk that can be inferred noisily from the project payoff. The mechanism design relies only on the reported risk aversion, as the equilibrium project risk choice depends on the reported risk aversion rather than the true risk aversion per se. The reason is that the manager's risk choice does not affect the mean payoff, so in the absence of moral hazard, shareholders do not condition the manager's wage on the project payoff. Thus, the manager's project risk choice is not driven by wage risk considerations and is independent of his true risk aversion. However, the project payoff is used to infer managerial ability and overconfidence. These inferences are used in conjunction with the manager's reported risk aversion in choosing a CEO at the end of the period. This causes the manager to base project risk choice on his reported risk aversion.

We assume that managers with revealed ability less than a cutoff are fired and suffer a personal cost due to firing.[20] The monetary equivalents of a manager's private benefit from promotion and personal cost of being fired are W^B and W^D, respectively.[21]

We seek to design a revelation mechanism M: $\Lambda \to \Re^2_+$ using linear wage contracts: $w_i = \beta(\sigma_i)Q + t(\sigma_i)$, where the piece rate $\beta \in \Re_+$ and the fixed wage $t \in \Re_+$ for manager i are functions of his reported risk aversion σ_i, and Q is an *exogenous* risky outcome with expected value of zero.[22]

The timing of events is as follows. First, the principal offers and precommits to a menu of wage contracts corresponding to types (risk aversion) of

[20] The firing of low ability managers is not critical to the model. It just ensures that managers do not take the maximum possible risk. This condition will also be satisfied if the promotion rule punishes excessive revealed overconfidence and thereby discourages excessive risk-taking.

[21] The manager's benefit from promotion and the cost of firing are expressed in monetary equivalents rather than utility terms for meaningful comparison across managers with different utility functions.

[22] An example of a real-world contract that bases the agent's compensation on an uncertain outcome that reveals no information ex post about the agent's action choice is executive compensation based on the firm's total return as opposed to idiosyncratic return, which filters out the market component. Prendergast (1999) cites the prevalence of contracts that expose managers to risky outcomes that are uninformative about managerial actions.

managers.[23] Then, all the managers simultaneously report their types and accept the corresponding wage contracts. Next, each manager chooses unobserved project risk. Finally, the risky outcome Q and project payoffs are realized. The outcome Q determines managerial wages, while project payoffs are used for promotion and firing decisions.

We examine this model backwards. The shareholders' choice of which manager to appoint as CEO in the second period and which managers to fire depends on the shareholders' posterior beliefs about managers' abilities, risk preferences, and overconfidence. Since these rules can be anticipated by the managers, all managers know how their contract choices and project risk choices will affect their chance of being promoted or getting fired. Next, we examine managers' project risk choices. When choosing project risk R_i, manager i trades off the probability of being promoted against the probability of being fired. While the utility from promotion and the disutility from dismissal do not vary across managers, the rules used to promote and fire a manager depend on the manager's reported risk aversion in the revelation game. Thus, manager i's project risk choice, a response to promotion and firing rules, will be determined by the risk aversion $\sigma_i^{'}$ reported by the manager. Consequently, the expected benefit from future promotion or the expected cost of future dismissal for manager i depends on the reported risk aversion $\sigma_i^{'}$.

Finally, in the wage contract choice stage, each manager reports a risk aversion to maximize his expected utility from his wage as well as from potential promotion and dismissal. We simplify by assuming that $t(\sigma_i^{'})$ includes the fixed part of manager i's wage and the monetary equivalent of future promotion and dismissal, and it depends on the manager's report $\sigma_i^{'}$. A sufficient condition for an implementable mechanism that elicits truthful reporting of managerial risk aversion, that is, $\sigma_i^{'} = \sigma_i$, is the Spence–Mirrlees single-crossing condition.[24] The single-crossing condition holds if

$$\frac{\partial}{\partial \sigma_i}\left(\frac{\partial E u_{\sigma_i} / \partial \beta}{\partial E u_{\sigma_i} / \partial t}\right)$$

is monotonic in σ_i, where the expectation Eu_{σ_i} is based on manager i's information and beliefs at the time he reports σ_i, and the manager takes as given the probability distribution of types and the actions of the other managers. This condition requires that a manager's marginal rate of substitution between his fixed wage and his payoff-contingent wage is monotonic in the manager's risk aversion.

PROPOSITION 7: *Shareholders can offer a menu of wage contracts $(\beta(\sigma'), t(\sigma'))$ with $\beta\left(\sigma'\right)$ decreasing in the reported risk aversion σ' such that each manager will truthfully report his risk aversion.*

The above result verifies that the single-crossing condition holds in our model and establishes that shareholders can implement a perfectly separating,

[23] We ignore managers' participation constraints because their participation can be ensured with sufficiently high fixed wages.

[24] See, for example, Fudenberg and Tirole (1991, p. 259).

risk-aversion revealing mechanism. The managers' choices are revealing because more risk-averse managers prefer contracts with lesser payoff dependence, even with lower expected wages. Shareholders can use their knowledge about a manager's risk aversion to infer the manager's risk choice and hence sharpen their inference of managerial ability from the project payoff. Further, this information affects the optimal promotion rule as well as wage contracts in the second period.[25] An analogous revelation mechanism, however, cannot be used to learn the degree of overconfidence because managers are not aware of their own overconfidence.

Our result that overconfident agents can be screened on the basis of risk aversion is in contrast to the existing literature. For example, Sandroni and Squintani (2007) find that overconfidence frustrates the separation of low-risk and high-risk agents by contract choice in compulsory insurance markets. In their model, a high-risk overconfident agent mistakenly believes he is low risk and underestimates the risk that he potentially seeks to insure. The consequent lack of behavioral separation between those who are truly low risk and those who are high risk but mistakenly believe they are low risk makes the usual Rothschild and Stiglitz (1976) separation of types by insurable risks impossible.

The reason we are able to achieve separation despite overconfidence is a key difference in the two models. Whereas in Sandroni and Squintani (2007) agents vary in overconfidence and the magnitude of the unidimensional risk they face, in our model agents vary in overconfidence and risk aversion. Two features of our model enable separation by risk aversion. The first is the recognition that while risk aversion pertains to an agent's consumption from *all* sources, an agent's overconfidence is *specific* to his private information.[26] Second, we permit the feasible contracting space in the design of the revelation mechanism to include an exogenous risky outcome unrelated to the manager's private information or domain of personal association. That is, we assume that there is a source of risk on which the principal and the agent can contract and over which the agent is risk averse but not overconfident. This is plausible, given the abundance of evidence in psychology that personal involvement heightens overconfidence without affecting risk tolerance with respect to exogenous risky outcomes. For example, Wu and Knott (2006) present evidence that entrepreneurs are risk averse with respect to (exogenous) demand uncertainty and overconfident with respect to self-ability uncertainty. Concrete examples of exogenous risky outcomes over which contracting can occur are the overall stock market return, the economic growth rate, and the weather (which can affect the demand for, say, air conditioners).

E. Effect of Overconfidence on Information Acquisition

The efficiency of the CEO's project portfolio investment decision depends on the information signal s available to her. So far, we have taken the precision of

[25] While we have shown that a risk aversion-revealing mechanism can be implemented, we have not addressed the issue of the efficiency of the mechanism and whether the firm would elicit information about risk aversion this way.

[26] This is consistent with the evidence on overconfidence (see Hirshleifer (2001)).

s as exogenous. We now allow the CEO to determine this precision and consider how her choice affects firm value.

Suppose after the CEO discovers the signal precision $q(C)$ based on her possibly biased beliefs, she can invest in information to improve this precision. An increase in precision from q_1 to $q_2 > q_1$ results in an information acquisition cost of $\theta(q_2) - \theta(q_1)$ to shareholders, where $\theta' > 0$, $\theta'' > 0$. This investment in information may eventually be observed by the board but is noncontractible. However, we ignore any agency problem in the CEO's information acquisition decision, so the CEO's precision choice seeks to maximize firm value net of information acquisition costs.

LEMMA 3: *The shareholders' expected payoff from the project portfolio is increasing in q, the precision of the project portfolio quality signal. An overconfident CEO underinvests in information acquisition.*

The first result in the lemma is intuitive. The investment decision of the CEO depends on portfolio quality as determined by the signal. This investment decision is like a real option with a convex payoff because the CEO invests in higher quality (higher payoff) portfolios and rejects lower quality portfolios. A more informative signal causes greater quality differentiation between portfolios, resulting in more efficient investment decisions. The intuition for the second result is as follows. The optimal investment in information depends on the initial precision of the CEO's signal and the desired precision of the signal. The desired precision is determined by the tradeoff between the marginal increase in the shareholders' payoff from higher precision and the marginal cost of increasing precision. An overconfident CEO desires the same precision as the rational CEO, but believes her original signal to be more precise and so believes the desired precision is attainable with a lower investment in information.[27] Thus, the overconfident CEO produces less precise information.[28]

[27] A robustness issue here is whether the CEO's overconfidence would result in overinvestment in information if the overconfidence was about her skill in producing information. This is, of course, possible, and it may be an interesting avenue to explore in the context of managers whose expertise is in information production/acquisition, such as those in charge of financial intermediaries (e.g., Ramakrishnan and Thakor (1984)). However, our focus here is on managers for whom information production is secondary to the task of choosing projects, and thus we focus on their overconfidence about the signal pertaining to this project choice rather than their skill in acquiring information. Our analysis is more in line with papers such as Daniel, Hirshleifer, and Subrahmanyam (1998) where the manager receives a signal whose precision he overestimates. A difference is that we permit the CEO's information acquisition choice to exert an independent influence on signal precision. For a paper that shows how overconfidence can lead to overinvestment in information, see Gervais, Heaton, and Odean (2007).

[28] Milbourn, Shockley, and Thakor (2001) show that career concerns may induce rational managers to *overinvest* in information production. This suggests that the underinvestment in information caused by CEO overconfidence may help mitigate the overinvestment in information caused by the CEO's career concerns.

PROPOSITION 8: *Firm value is higher with a moderately overconfident CEO than with a rational CEO when the CEO is responsible for both project portfolio investment and information acquisition decisions.*

The above result shows that moderate overconfidence in the CEO is good for the shareholders. Proposition 5 showed that moderate overconfidence can mitigate underinvestment, and Lemma 3 showed that overconfidence leads to suboptimally low information acquisition and hence less efficient investment decisions. Thus, when both investment efficiency and information acquisition efficiency are considered, we see that CEO overconfidence pulls in opposite directions—it enhances investment efficiency but diminishes information acquisition efficiency. This makes the desirability of CEO overconfidence potentially ambiguous for shareholders. Proposition 8 indicates, however, that for relatively low levels of overconfidence, the inefficiency associated with underinvestment in information is more than offset by the gain from the reduction in underinvestment in the project portfolio due to CEO risk aversion.[29]

V. CEO Overconfidence and Corporate Governance at the Board Level

The purpose of this section is to examine board governance. Instead of assuming that the first-period CEO is pre-determined to retire at the end of the first period, we now permit the board to decide whether to retain or fire her based on the first-period outcome. The structure of the model is similar to that in Section IV in that we suppress individual managers' risk choices and abilities and focus on the strategy choice and ability of the CEO. The two differences are that now we fix the precision of the signal s, suppressing the CEO's ability to affect it, and we allow the CEO's first-period decision-making to recognize that her decision in the first period will affect not only her compensation but also the board's decision to retain/fire her after the first period. That is, unlike the "end game" in Section IV, there are now career concerns. This relates to the first-period analysis, which is our focus here. We don't analyze the second period since all our results from the previous section carry over intact to the second period.

A. *The Board's Decision to Retain/Fire the CEO*

The board can observe the firm's payoff and whether the CEO accepted the project portfolio. In addition, the board acquires a signal about the quality of the accepted portfolio; this signal is non-contractible. For simplicity, we assume

[29] Of course, what level of overconfidence is considered "relatively low" depends on risk aversion. The higher the CEO's risk aversion, the higher will be the level of overconfidence that is considered low.

that the board's signal is the same as the signal s that the CEO observes.[30] The signal permits the board to improve its evaluation of the CEO's investment decision since the board can now observe the information the CEO had at the time of the investment decision.

The board forms beliefs about the project portfolio and updates its beliefs about the CEO's ability and overconfidence. The updated beliefs are used to decide whether the CEO is retained for the second period. If retained, the CEO gets a new wage contract for the second period that depends on the board's updated beliefs about the CEO. If the CEO is fired, the replacement CEO's ability and overconfidence are unknown and drawn from distributions that characterize the population of potential replacement CEOs.

The board's CEO-replacement decision in the first period maximizes the expected second-period payoff of the firm net of the (new) CEO's wage. The board's beliefs about the abilities and overconfidence of the incumbent CEO and the replacement CEO may differ; these beliefs affect the reservation wages, the optimal wage contracts, and consequently the expected firm payoffs across the two CEOs.

Let ψ be the probability density function describing the board's prior belief about the CEO's ability A_0 at the beginning of the first period and let ζ be the corresponding posterior distribution at the end of the first period. Let μ be the probability density function describing the board's prior belief about the CEO's degree of overconfidence C at the beginning of the first period and let ν be the corresponding posterior distribution at the end of the first period.

Lemma 2 and Proposition 5 imply that the firm's expected second-period payoff is increasing in CEO ability but nonmonotonic in the CEO's overconfidence. In our analysis of the first-period problem, using subscripts to indicate time, we can therefore assume that the CEO expects a reservation utility of $U_2(\zeta, \nu)$ and the shareholders expect a payoff of $Y_2(\zeta, \nu)$ in the second period, where U_2 and Y_2 depend on the posterior distributions $\zeta(A_0)$ and $\nu(C)$ of the CEO's ability and overconfidence as follows:

$$U_2(\zeta, \upsilon) = \int \int U_2(A_0, C)\nu(C)\, dC \zeta(A_0)\, dA_0$$

and

$$Y_2(\zeta, \upsilon) = \int \int Y_2(A_0, C)\nu(C)\, dC \zeta(A_0)\, dA_0.$$

Here $U_2(A_0, C)$ and $Y_2(A_0, C)$ are linearly increasing in CEO ability A_0 and single-peaked in CEO overconfidence C in accordance with Lemma 2 and Proposition 5, respectively.

The board updates its beliefs about the CEO's ability and overconfidence based on the signal s, the outcome ω, and the project-specific payoffs X. We

[30] This is not necessary, and our results will also follow, albeit with more algebra, if the two signals are positively correlated.

represent the posterior distribution of the CEO's overconfidence as $\nu(C \mid s, \omega, X)$, and the posterior distribution of the CEO's ability as $\zeta(A_0 \mid s, \omega, X)$.[31]

The CEO now makes her portfolio choice in the first period to maximize her expected utility over two periods, taking into account the impact of her decision on the probability of being retained for the second period. That is, the CEO maximizes

$$E[U(W) + U_2(\zeta, \nu)] + [B \times \Pr(Y_2(\zeta, \nu) > Y_2(\psi, \mu))]. \tag{15}$$

Note that we have assigned the same private benefit B to the CEO from being retained in the second period that we included in the utility specification of the manager who gets promoted to be CEO. The CEO's reservation utility is U_1. The board chooses wages W_H, W_L, and W_R and the CEO determines the investment policy. We restrict attention to equilibria in which if the CEO accepts a portfolio with a lower signal value, she will also accept a portfolio with a higher signal value. This requires $W_H \geq W_L$. Further, if a less overconfident CEO accepts a portfolio based on a signal s, so will a more overconfident CEO who overreacts to the good news in the signal. Thus, a CEO with overconfidence $\hat{C}$ will invest when the portfolio quality signal exceeds $s^*(\hat{C})$, where s^* is a decreasing function. The board's inference from observing the CEO accept (reject) a portfolio with signal s will be that the signal s exceeds (is less than) the CEO's threshold $s^*(C)$. The board's problem is

$$\underset{s^*(C), W_H, W_L, W_R}{\text{Max}} \left\{ \begin{array}{l} E\left[\sum_{i=1}^{n} x_i\right] \\ + \Pr(s \geq s^*(C)) \times E\left[p(s,1)(h - W_H) + \{1 - p(s,1)\}(l - W_L) \mid s \geq s^*(C)\right] \\ + \Pr(s < s^*(C)) \times (r - W_R) \\ + E[\max(Y_2(\psi, \mu), Y_2(\zeta, \nu))] \end{array} \right\} \tag{16}$$

subject to

$$u_A(s^*(\hat{C}), \hat{C}) = u_R(s^*(\hat{C}), \hat{C}) \;\; \forall \hat{C}, \tag{17}$$

$$W_H \geq W_L, \tag{18}$$

$$u_R(s, \hat{C}) \geq u_L(s, \hat{C}) \;\; \forall s, \hat{C}, \tag{19}$$

and

$$\begin{aligned} &\Pr(s \geq s^*(C)) E\left\{u_A(s, C) \mid s \geq s^*(C)\right\} \\ &\quad + \Pr(s < s^*(C)) E\left\{u_R(s, C) \mid s < s^*(C)\right\} \geq U_1, \end{aligned} \tag{20}$$

[31] The results in Lemma 2 and Propositions 5 and 8 hold the CEO's reservation utility fixed. Common knowledge about the CEO's ability and overconfidence may affect this reservation utility. We assume that even when reservation utility changes, the changes are small compared to the changes in firms' expected payoffs, so that the results in Lemma 2 and Propositions 5 and 8 continue to hold for firm value net of CEO wages.

where

$$u_H(s,\hat{C}) = u(W_H) + E\left\{U_2(\zeta,\nu) + B \times \Pr(Y_2(\zeta,\nu) > Y_2(\psi,\mu))\,|\,s \geq s^*(C)\right\}, \quad (21)$$

$$u_L(s,\hat{C}) = u(W_L) + E\left\{U_2(\zeta,\nu) + B \times \Pr(Y_2(\zeta,\nu) > Y_2(\psi,\mu))\,|\,s \geq s^*(C)\right\}, \quad (22)$$

$$u_R(s,\hat{C}) = u(W_R) + E\left\{U_2(\zeta,\nu) + B \times \Pr(Y_2(\zeta,\nu) > Y_2(\psi,\mu))\,|\,s < s^*(C)\right\}, \quad (23)$$

and

$$u_A(s,\hat{C}) = p(s,\hat{C})u_H(s,\hat{C}) + \{1 - p(s,\hat{C})\}u_L(s,\hat{C}) - c. \quad (24)$$

The objective in (16) is the expected payoff to shareholders from both periods. It is similar to (7) except for the addition of a fourth term, $E[\max(Y_2(\psi,\,\mu), Y_2(\zeta,\,\nu))]$, which represents the expected value of the maximum of two second-period payoffs, one from retaining the current CEO and the other from hiring a new CEO. Equation (17) defines the threshold signal value $s^*(\hat{C})$ for portfolio acceptance. Constraint (18) ensures that the CEO will accept a portfolio when the signal value exceeds $s^*(\hat{C})$, because acceptance increases the probability of obtaining the higher wage W_H and precludes the shareholders from considering the CEO diffident. Constraint (19) ensures that the CEO prefers rejecting a portfolio to accepting it but not developing it. The Individual Rationality (IR) constraint (20) ensures the CEO's participation.

The expected utilities in (21), (22), (23), and (24) are based on the potentially biased beliefs of the CEO with overconfidence $\hat{C}$ and signal s, with the respective outcomes of a good portfolio, a bad portfolio, portfolio rejection, and portfolio acceptance with the development cost. The CEO's expected utility is calculated based on her expectations of her first-period wage, the expected utility from her second-period wage, and the probability of being retained as CEO.

LEMMA 4: *If the deviation in the CEO's confidence from rationality is sufficiently small, there is an equilibrium in which:*

- *(a) A CEO with sufficiently low confidence underinvests: When the CEO is indifferent between accepting and rejecting a portfolio, the shareholders would be strictly better off if the portfolio is accepted.*
- *(b) Firm value is increasing in the CEO's ability A_0.*
- *(c) The expected value of the firm is increasing in the CEO's overconfidence for degrees of overconfidence less than a cutoff.*

This lemma shows that the results in Section IV for the second period continue to hold in the first period under certain conditions even when the CEO has career concerns. Career concerns may strengthen or weaken underinvestment depending on whether the CEO thinks that indicating greater confidence by accepting a portfolio will raise or lower her second-period utility. To the extent that the learning and career concerns are not very significant, CEO risk aversion towards her first-period wage still leads to underinvestment, at least if she is not too overconfident. A sufficiently overconfident CEO may overestimate portfolio quality so much that she overinvests, whereas moderate CEO overconfidence raises firm value by countering underinvestment.

LEMMA 5: *There is no learning about the CEO's ability from the project portfolio quality signal s or the portfolio outcome ($\omega = H$, L, or R). The posterior distribution ζ of the incumbent CEO's ability is increasing in each project-specific payoff x_i in the first-order-stochastic dominance sense. The posterior distribution ν of the incumbent CEO's degree of overconfidence is independent of X, the vector of project-specific payoffs, and depends only on ω and s. Conditional on portfolio acceptance or rejection, the distribution ν with a lower value of s first-order-stochastically dominates the distribution with a higher value of s. The posterior distribution ν when the portfolio is accepted first-order-stochastically dominates the prior distribution, which in turn first-order-stochastically dominates the posterior distribution ν when the portfolio is rejected.*

The intuition is that CEO ability affects only the distribution of the project-specific payoffs X; since X is observable, the signal s and the portfolio outcome ω do not affect beliefs about the CEO's ability. Higher values of x_i indicate higher CEO ability because the distribution of x_i with a higher-ability CEO first-order-stochastically dominates that of a lower ability CEO.

Learning about CEO overconfidence depends only on the portfolio quality signal s and the portfolio outcome ω—whether the portfolio was accepted and if so, whether it turned out to be good or bad. This learning is independent of project-specific payoffs X because overconfidence affects only the CEO's investment decision and has no impact on X. Since a more overconfident CEO uses a lower cutoff value of s to accept a portfolio, a lower s value for an *accepted* portfolio indicates higher CEO overconfidence, and a higher s value for a *rejected* portfolio indicates lower CEO overconfidence. Portfolio rejection decreases *perceived* CEO overconfidence because a portfolio is rejected only if CEO overconfidence is less than the value at which the CEO is indifferent between portfolio acceptance and rejection; likewise, portfolio acceptance raises perceived CEO overconfidence.

The decision of the board to retain or replace the incumbent CEO will be based on the learning about the CEO's ability and overconfidence from the first period. The board will compare the expected firm value with the incumbent CEO to the expected firm value with a new CEO.[32] The following result characterizes the conditions under which the incumbent CEO will be fired.

PROPOSITION 9: *For any portfolio quality signal value s and portfolio outcome ω, the CEO is fired if the project-specific payoffs X are sufficiently low and retained otherwise. For any X, the CEO is fired if she rejects the portfolio and s is sufficiently high, or if she accepts the portfolio and s is sufficiently low.*

[32] Hermalin (2005) shows that the option of firing the CEO in the future biases the board towards a candidate whose ability is known less precisely.

The intuition is as follows. The board's objective is to maximize shareholders' expected second-period payoff Y_2, which is affected by the CEO in three ways. First, the CEO's overconfidence will influence her decision to accept a portfolio. Second, an accepted portfolio's outcome ω will depend on whether the CEO developed the portfolio. Third, the CEO's ability will affect project-specific payoffs X. The equilibrium second-period contract will induce the CEO to develop an accepted portfolio, so the board focuses on the CEO's overconfidence and ability, which will influence investment policy and X, respectively, in the second period and thereby determine Y_2. Since Y_2 is increasing in the CEO's ability and the CEO's *inferred* ability is increasing in the first-period project-specific payoffs X, for given beliefs about CEO overconfidence based on s and ω, the CEO will be retained if project-specific payoffs X are sufficiently high and replaced otherwise. The expected second-period payoff Y_2 is single-peaked in CEO overconfidence, so conditional on beliefs about CEO ability, the CEO will be replaced if she is evaluated to be excessively diffident or excessively overconfident. If the CEO accepts a portfolio with a very high s or rejects one with a very low s, there is little learning about CEO overconfidence because most CEOs make these decisions. It is when the CEO rejects a portfolio for a sufficiently high s that it is possible to make the sharp inference that the CEO is excessively diffident. Similarly, when the CEO accepts a portfolio for a sufficiently low s, it very likely that she is excessively overconfident. Thus, overconfidence is more likely to precipitate CEO dismissal when the CEO overinvests, while diffidence is more likely to precipitate dismissal when the CEO underinvests.

B. The Impact of Information Disclosure-Related Penalties on the Board's Choice of CEO

The Sarbox specifies stiff penalties for misrepresentation of information by the firm. To examine the effect of this Act, we now add three features to the model. First, we assume that the signal s that the CEO uses to make her portfolio-specific decision at $t = 0$ and the board uses to decide whether to retain the CEO at $t = 1$ is also made available to investors at $t = 0$, even though the precision of this signal is known privately only to the CEO who makes an unobservable investment in this precision.

Second, if the portfolio payoff observed at $t = 1$ is sufficiently low, it is inferred that the precision of the information provided to investors at $t = 0$ was low enough to invoke disclosure penalties. This is because a low payoff indicates a high likelihood that the signal s was erroneous in indicating that the portfolio should be undertaken. We shall take this cutoff portfolio payoff, α^*, as exogenously given, so the penalty, $P \geq 0$, a reduction in firm value, is invoked if $\alpha < \alpha^*$.

Third, at date $t = -1$, when the board is deciding who to appoint as CEO at $t = 0$, the board knows that managers of equal assessed ability include those with confidence $C > 1$ (overconfident managers) and confidence $C = 1$ (rational managers). These managers cannot be noiselessly distinguished, but there are two observationally distinct groups, A and B, such that the probability that a randomly chosen manager will have $C > 1$ is higher for group A than for group B. We now have:

COROLLARY 2: *Suppose the CEO can choose the precision of the signal s at* $t = 0$ *by making an unobservable investment in information production. Then there exist parameter values such that the board prefers to choose a manager from group A as the incumbent CEO at* $t = -1$ *when* $P = 0$. *Moreover, there exists a* $P > 0$ *high enough such that the board's preference switches to a manager from group B as the incumbent CEO chosen at* $t = -1$.

This result points out that in the absence of information disclosure penalties, the board will prefer a moderately overconfident ($C > 1$ but not too large) CEO. The reason is that such a CEO attenuates the risk aversion-induced underinvestment inefficiency associated with a rational CEO, and this more than compensates for the inefficiencies of a less precise signal both in terms of the CEO's portfolio selection decision and the board's decision of whether to retain the incumbent CEO at $t = 1$. And since C is not too high, the overinvestment inefficiency with excessive overconfidence can be avoided.

An information imprecision penalty, P, can change this preference on the board's part. When P is sufficiently high and is a charge against firm value,[33] the board may find the cost of underinvestment in information precision to be too high to appoint a CEO who has a high likelihood of being overconfident. Even though the anticipation of this penalty increases the CEO's investment in information precision, an overconfident CEO always invests less in information precision than a rational CEO.

Corollary 2 highlights an unintended consequence of Sarbox. It can tilt CEO selection in favor of less confident CEOs, thereby increasing the incidence of *underinvestment* in projects.[34]

VI. Empirical Predictions

We now discuss the empirical predictions of the analysis and the results to which they are linked.

First, the incidence of overconfidence among CEOs will be higher than in the general population. This follows from Propositions 2 and 3. Moreover, our analysis indicates where CEO overconfidence is likely to be especially pervasive. Specifically, overconfident CEOs are more likely to be found in firms

[33] The charge would include both direct monetary penalties and the indirect effect of a reduction in firm value because litigation distracts management and the board.

[34] To the extent that Sarbox causes the board to eschew a moderately overconfident CEO who would have enhanced firm value in the absence of the Sarbox penalties, Sarbox will also lead to a reduction in firm values. However, this does not take into account potential benefits of Sarbox that are outside our model, such as improved investor confidence, reduction in self-serving behavior by managers, etc.

in riskier industries (see the discussion following Proposition 5) and in firms that emphasize merit-based promotions more (see the discussion following Proposition 3).

Second, from Proposition 5 and Lemma 4 we have the prediction that, controlling for industry, firm size, and future growth, the relationship between firm value and the CEO's overconfidence is non-monotonic, with firm value increasing with CEO overconfidence up to a point (see also Proposition 8) and then decreasing. The threshold overconfidence value above which CEO overconfidence diminishes firm value obviously depends on the CEO's risk aversion. However, it may be empirically impractical to construct continuous measures of risk aversion and overconfidence, and a null hypothesis of a "non-monotonic effect" may not be an ideal test. To get around both these issues, one could identify a sample of only low risk aversion managers. For this sample, our theory generates a null hypothesis with a clear direction: Overconfidence (even a rough binary measure) should reduce value.

Third, Proposition 5 also predicts stock price reactions to corporate investments. To see this, note first that our model assumes that the CEO's overconfidence is unknown to everyone, although it is common knowledge that there is a nonzero probability that any manager is overconfident. Once a CEO has been at the helm for some time, we would expect that the stock market has formed a noisy yet informative assessment of the CEO's overconfidence—perhaps using proxies such as those used by Malmendier and Tate (2005)—so that the market's posterior assessment of CEO overconfidence may rise above its prior assessment even though the CEO's own prior belief about her overconfidence may be unaffected.[35] The market will thus end up with cross-sectional heterogeneity in its posterior beliefs about the degrees of overconfidence of various CEOs. This leads to the prediction that the (abnormal) price reaction to a corporate investment that has an identifiable announcement date will be higher for CEOs viewed by the market as possessing relatively low to moderate levels of overconfidence than for those viewed by the market as being highly overconfident. This prediction too is really about the *interaction* between overconfidence and risk aversion, so one would want to control for CEO risk aversion and test whether the predicted relationship between overconfidence and the announcement effect holds.

Fourth, from Lemma 3 we know that an overconfident CEO underinvests in information, which implies that the quality of the information provided to both the board and the investors will be poorer with more overconfident CEOs.

[35] With attribution bias (see Hirshleifer (2001)) feeding the CEO's overconfidence, it follows that even though the CEO and the market observe the same realizations of events, the market may revise upward its belief about the CEO's overconfidence, whereas the CEO's belief either does not change or changes in the opposite direction. We could even permit the board to share the stock market's belief about the CEO's overconfidence. However, as our analysis in Section V shows, this does not guarantee that the board will dismiss the CEO. There are two reasons for this. First, the board views moderate levels of CEO overconfidence as being beneficial to the shareholders. Second, even for high levels of overconfidence that do not benefit the shareholders, the board may not learn enough to fire the CEO unless the overconfidence is excessive (see Section V).

One way to test this prediction would be to identify all settled litigation cases involving information misrepresentation to investors, and examine whether the CEOs in this group were more overconfident than for a control sample, making sure to introduce appropriate control variables.

Fifth, Proposition 9 yields the prediction that, controlling for CEO risk aversion, CEOs viewed as either excessively diffident or excessively overconfident are fired. To test this prediction, one would examine the attributes of dismissed CEOs and compare those to the attributes of CEOs who have been in office for a while. Controlling for risk aversion, there should be a greater clustering of dismissed CEOs in the two tails of the overconfidence distribution and a predominance of continuing CEOs in the middle.

Sixth, Proposition 9 also generates a predicted announcement effect. There will be a positive (abnormal) stock price reaction to a CEO dismissal, and the size of the reaction will be increasing in the absolute value of the difference between the (appropriately normalized) level of pre-dismissal investment by the firm and the contemporaneous average industry investment, controlling for other factors that impact corporate investment. The idea is that, while on average, CEOs may be overconfident, the level of overconfidence among retained industry CEOs is closer to the shareholders' optimum than that of a dismissed CEO. An excessively overconfident CEO is fired for overinvesting and his overinvestment is increasing in his overconfidence, which implies that the stock price reaction will be increasing in the amount of overinvestment. We would look at the absolute value of the difference between the firm's investment and the average industry investment because diffident CEOs who underinvest also get fired.

Finally, from Corollary 2 we get the prediction that Sarbox should lead to a lower incidence of overconfidence among CEOs, and lower aggregate investment. A test of this prediction should be possible to conduct now, as there should be sufficient CEO turnover since the passage of Sarbox to generate a post-Sarbox sample of CEOs that differs potentially significantly from the pre-Sarbox sample.

VII. Conclusion

We examine how corporate governance at the board level is affected by the internal organizational governance that seeks to identify the highest ability manager to appoint as CEO through the ability filtering provided by an implicit tournament. We find that the board is likely to end up with a pool of overconfident managers from which to choose a CEO and that moderate degrees of overconfidence actually benefit the shareholders. However, excessively overconfident CEOs overinvest, and firm value is nonmonotonic in CEO overconfidence. Hence, a board acting in the shareholders' best interest will fire a CEO who is either perceived to be overly diffident or overly overconfident. That is, internal organizational governance induces the board to appoint an overconfident manager as CEO, but the board fires the CEO if it is later discovered that she is too overconfident. Our theory explains the apparent paradox that overconfident CEOs sometimes make value-destroying investments, and yet overconfident managers are winners in the race to be CEO. In our analysis, we have taken seriously the important task of theoretically distinguishing overconfidence from risk aversion. In addition to these main results, our analysis generates numerous testable predictions that are summarized above.

We believe that the interaction between CEO overconfidence and the effectiveness of corporate governance is an important issue. Future research involving CEO attributes associated with endogenous selection processes should lead to a more comprehensive framework for assessing corporate governance.

Appendix

Proof of Lemma 1: Suppose manager 1 chooses risk R while all other managers choose risk $\hat{R} < R$. The probability that manager 1 is promoted is

$$P_1 = \Pr(x_1 > x_j \forall j > 1) = \int_{-\infty}^{\infty} \{1 - F(x, R)\} \bar{f}(x)\, dx,$$

where $\bar{f}$ is the probability density function of the maximum of $x_2, \ldots, x_n$. Rearranging the above, we get

$$\begin{aligned} P_1 &= \int_{-\infty}^{\infty} \{1 - F(x, \hat{R})\} \bar{f}(x)\, dx + \int_{-\infty}^{\infty} \{F(x, \hat{R}) - F(x, R)\} \bar{f}(x)\, dx \\ &= \frac{1}{n} + \int_{0}^{\infty} \{F(x, \hat{R}) - F(x, R)\}\{\bar{f}(x) - \bar{f}(-x)\}\, dx \geq \frac{1}{n}, \end{aligned} \tag{A1}$$

where the first term of $1/n$ after the equality is the probability a manager will be promoted when all managers choose the same risk $\hat{R}$ and the second term is obtained from the property $F(-x, \hat{R}) - F(-x, R) = 1 - F(x, \hat{R}) - 1 + F(x, R) = -\{F(x, \hat{R}) - F(x, R)\}$. The inequality obtains because with $x > 0$, $F(x, \hat{R}) - F(x, R) > 0$ and $\bar{f}(x) = \bar{f}(-x)$ for $n = 2$, while if $n > 2$,

$$\begin{aligned} \bar{f}(x) &= \frac{d}{dx}\{F(x)\}^{n-1} = (n-1)\{F(x)\}^{n-2} f(x) \\ &> (n-1)\{F(-x)\}^{n-2} f(-x) = \bar{f}(-x). \end{aligned}$$

Q.E.D.

Proof of Proposition 1: Let $E^L[u_1 \mid R, R^b]$ be the expected utility of manager 1 when the benefit from promotion is $B_L \geq 0$, and $E^H[u_1 \mid R, R^b]$ be the corresponding expected utility when the benefit from promotion is $B_H > B_L$ under the assumptions that shareholders expect all managers to choose risk R^b, manager 1 chooses risk R, and the other managers choose R^b. The difference in the two utilities is given by

$$E^H[u_1 | R, R^b] - E^L[u_1 | R, R^b] = (B_H - B_L) \int_{-\infty}^{\infty} \{1 - F(x, R)\} \bar{f}(x)\, dx,$$

where $\bar{f}$ is defined in the proof of Lemma 1. Differentiating the above expression with respect to R and evaluating at $R = R^b$, we get

$$\begin{aligned}
&\frac{d}{dR}E^H[u_1|R, R^b]|_{R=R^b} - \frac{d}{dR}E^L[u_1|R, R^b]|_{R=R^b} \\
&= -(B_H - B_L)\int_{-\infty}^{\infty}\frac{dF(x, R^b)}{dR^b}\bar{f}(x)\,dx \qquad \text{(A2)} \\
&= -(B_H - B_L)\int_{0}^{\infty}\frac{dF(x, R^b)}{dR^b}\{\bar{f}(x) - \bar{f}(-x)\}\,dx > 0,
\end{aligned}$$

where the last equality and the last inequality follow from the arguments in the proof of Lemma 1 and the fact that $dF(x, R^b)/dR^b < 0$ for $x > 0$. Suppose $R^b < R^{\max}$ is the equilibrium risk choice with benefit from promotion B_L. Then, with benefit from promotion B_H, when other managers choose risk R^b the best response of manager 1 is a risk choice higher than R^b, while when the other managers choose risk $R^{\max}$ the best response of manager 1 is a lower risk choice. By continuity, there is an equilibrium risk level between R^b and $R^{\max}$. The case without promotion concerns is equivalent to $B_L = 0$. Q.E.D.

Proof of Proposition 2: Follows from Lemma 1 by assuming manager 1 is the overconfident manager and chooses risk $R = C\hat{R} > \hat{R}$ when all other managers choose risk $\hat{R}$. Manager 1's promotion probability is increasing in the degree of overconfidence C because an increase in C increases the risk R of manager 1's project, which increases $F(x, \hat{R}) - F(x, R)$ for $x > 0$, thereby increasing the value of the integral in (A1). Q.E.D.

Proof of Proposition 3: Since the managers do not know whether they are rational or overconfident and are otherwise identical, they face the same project risk choice problem and each manager chooses the same project risk R^{**}, knowing that the actual risk will be CR^{**} if he is overconfident. Since an overconfident manager's project is riskier than a rational manager's project, it follows from Lemma 1 that an overconfident CEO is more likely to be promoted than a rational manager. Q.E.D.

Proof of Proposition 4: Constraint (9) can be eliminated as it follows from (8) and (10). The board's problem then reduces to

$$\underset{p^*, W_H, W_L, W_R}{\text{Max}} \left\{ \begin{array}{l} E\left[\sum x_i\right] \\ + \Pr(p(s, C) \geq p^*) \times E\left[p(s, 1)(h - W_H) + \{1 - p(s, 1)\}(l - W_L) \mid p(s, C) \geq p^*\right] \\ + \Pr(p(s, C) < p^*) \times (r - W_R) \end{array} \right\} \tag{A3}$$

subject to

$$p^* u(W_H) + (1 - p^*)u(W_L) - c = u(W_R), \tag{A4}$$

$$W_R \geq W_L, \tag{A5}$$

and

$$\Pr(p(s,C) \geq p^*) \times E\left[p(s,C)u(W_H) + \left\{1 - p(s,C)\right\} u(W_L) - c \mid p(s,C) \geq p^*\right] + \Pr(p(s,C) < p^*) \times u(W_R) \geq u(W_0). \tag{A6}$$

Let η, $\kappa \leq 0$, and $\delta \leq 0$ be the Lagrange multipliers for constraints (A4), (A5), and (A6), respectively, in the above problem. The Individual Rationality (IR) constraint (A6) must bind because otherwise the expected CEO wage can be reduced without violating other constraints by lowering $u(W_H)$, $u(W_L)$, and $u(W_R)$ by the same constant. Next, to prove $\eta < 0$, consider the optimal solution in absence of constraint (A4). If (A5) binds, the solution is

$$\begin{aligned} u'(W_R) &= u'(W_L) \\ &= \frac{E[p(s,C) \mid p(s,1) \geq p^*]}{E[p(s,1) \mid p(s,1) \geq p^*]} \frac{1 - \Pr(p(s,1) \geq p^*)E[p(s,1) \mid p(s,1) \geq p^*]}{1 - \Pr(p(s,1) \geq p^*)E[p(s,C) \mid p(s,1) \geq p^*]} u'(W_H). \end{aligned}$$

If (A5) does not bind, the solution to (A3) and (A6) is

$$\begin{aligned} u'(W_H) &= \frac{E[p(s,1) \mid p(s,1) \geq p^*]}{E[p(s,C) \mid p(s,1) \geq p^*]} u'(W_R), \\ u'(W_L) &= \frac{1 - E[p(s,1) \mid p(s,1) \geq p^*]}{1 - E[p(s,C) \mid p(s,1) \geq p^*]} u'(W_R). \end{aligned}$$

In both cases, as long as the CEO's expected deviation from rationality ($C = 1$) is small, W_H and W_L are close to W_R so the right-hand side of (A4) exceeds the left-hand side. Thus, we must have $\eta < 0$ in solution of (A3) to (A6). Next, let $\Gamma(p)$ be the expected payoff of the shareholders in (A3) with equilibrium wages conditional on threshold project quality p. The first-order condition for threshold portfolio quality to be p^* in equilibrium is

$$\begin{aligned} &\Gamma'(p^*) - \eta(u(W_H) - u(W_L)) \\ &\quad -\delta\left[u(W_R) + c - p^* u(W_H) - (1 - p^*)u(W_L)\right] \frac{d}{dp^*} \Pr\left[p(s,C) \leq p^*\right] = 0. \end{aligned}$$

Substituting $\eta < 0$, (9), and (A4), we get

$$\Gamma'(p^*) < 0.$$

Thus, the shareholders' expected payoff will increase if the CEO starts accepting portfolios with assessed quality marginally less than p^*. Q.E.D.

Proof of Lemma 2: Consider the optimal contract when the CEO's ability is known to be A_L. The same contract is feasible when the ability of the CEO is known to be $A_H > A_L$ because ability affects only the distribution of project-specific payoffs X, and the wage is not contingent on X. The expected wage of the CEO remains unchanged while the expected values of project-specific payoffs x_i increase because the distribution of x_i with CEO ability A_H

first-order-stochastically dominates the distribution of x_i with CEO ability A_L. The firm value increases. Q.E.D.

Proof of Proposition 5: The overconfident CEO accepts a portfolio if $s \geq s^C$ such that $p(s^C, C) = p^* = p(s^*, 1)$. That is,

$$s^C q(C) + 0.5[1 - q(C)] = s^* q^* + 0.5(1 - q^*). \tag{A7}$$

It is straightforward to check that

$$\mathrm{d}s^C/\mathrm{d}C < 0. \tag{A8}$$

Let $V(C)$ be the firm value when the CEO's degree of overconfidence is C:

$$V(C) = E\left[\sum x_i\right] + \int_0^{s^C} (r - W_R)\,ds + \int_{s^C}^1 \{p(s,1)[h - W_H] + [1 - p(s,1)](l - W_L)\}\,ds.$$

Differentiating with respect to C,

$$V'(C) = \frac{ds^C}{dC}\left\{(r - W_R) - p(s^C,1)[h - W_H] - [1 - p(s^C,1)](l - W_L)\right\}.$$

Since s^C is decreasing in C (from (A8)) and $p(s, 1)$ is increasing in s, we get $V'(C) > 0$ for $C < C^*$ and $V'(C) < 0$ for $C > C^*$, where at $C = C^*$,

$$p(s^C, 1) = p^{**} \equiv \frac{(r - W_R) - (l - W_L)}{[h - W_H] - (l - W_L)}.$$

But the value of s^C defined by the above equation is s^{**}. Substituting this in (A7) and simplifying, we characterize C^* by

$$q(C^*) = \frac{s^* - 0.5}{s^{**} - 0.5} q^* = \frac{p^* - 0.5}{p^{**} - 0.5} q^*. \qquad \text{Q.E.D.}$$

Proof of Proposition 6: The board's problem (A3) to (A6) with a rational CEO reduces to

$$\underset{p^*, W_H, W_L, W_R}{\mathrm{Max}} \left\{ \begin{array}{l} E\left[\sum x_i\right] + \Pr(p \geq p^*) \times E\left[p(h - W_H) + (1 - p)(l - W_L) \mid p \geq p^*\right] \\ + \Pr(p < p^*) \times (r - W_R) \end{array} \right\} \tag{A9}$$

such that

$$p^* u(W_H) + (1 - p^*) u(W_L) - c = u(W_R), \tag{A10}$$

$$\gamma \equiv u(W_R) - u(W_L) \geq 0, \tag{A11}$$

$$\Pr(p \geq p^*) \times E\left[pu(W_H) + (1-p)u(W_L) - c \mid p \geq p^*\right] \\ + \Pr(p < p^*) \times u(W_R) \geq u(W_0). \tag{A12}$$

The Individual Rationality constraint (A12) must bind because otherwise reducing $u(W_H)$, $u(W_L)$, and $u(W_R)$ by the same constant will not violate (A12), will leave the IC constraints (A10) and (A11) unchanged, and will increase the objective in (A9).

Taking p^* and γ as given, we can solve (A10), (A11), and (A12) for wages:

$$u(W_H) = u(W_0) + \frac{1-\tau}{p^*}c + \left(\frac{1-\tau}{p^*} - 1\right)\gamma, \tag{A13}$$

$$u(W_L) = u(W_0) - \frac{\tau}{p^*}c - \left(1 + \frac{\tau}{p^*}\right)\gamma, \tag{A14}$$

$$u(W_R) = u(W_0) - \frac{\tau}{p^*}c - \frac{\tau}{p^*}\gamma, \tag{A15}$$

and

$$\tau \equiv E[\max(p - p^*, 0)].$$

We now argue that γ must be zero. For $\gamma > 0$, $W_H > W_R > W_L$ and a decrease in γ reduces $W_H - W_R$ and $W_R - W_L$, reducing wage risk. Since the CEO is risk averse, this lowers expected CEO wage. Hence, the optimal solution must have $\gamma = 0$, that is, (A11) binds.

Now, consider a set of individuals as possible CEOs with utility functions u_σ indexed by σ, the risk aversion of the individual. The individuals are otherwise equivalent as CEOs. That is, the reservation wage for each is W_0 when they are paid a fixed wage and do not incur the cost of developing a portfolio, the reservation wage for each is W_1 when they are paid a fixed wage and incur the cost of developing a portfolio, and the distributions of their abilities are identical. Suppose the same investment policy is implemented for all CEOs. That is, p^* is fixed. From (A10) to (A12),

$$u_\sigma(W_H) = \frac{1-\tau}{p^*}u_\sigma(W_1) - \left(\frac{1-\tau}{p^*} - 1\right)u_\sigma(W_0), \tag{A16}$$

and

$$u_\sigma(W_L) = u_\sigma(W_R) = \left(1 + \frac{\tau}{p^*}\right)u_\sigma(W_0) - \frac{\tau}{p^*}u_\sigma(W_1). \tag{A17}$$

Since $W_1 > W_0$ and a higher σ increases the concavity of u_σ, application of Jensen's theorem to (A16) and (A17) shows that W_L, W_H, and W_R are increasing in σ. Since the expected CEO wage is a weighted average of W_L, W_H, and W_R

(see (A9)), a less risk-averse CEO implements the same investment policy with lower expected wages. Q.E.D.

Proof of Proposition 7: Manager i's expected utility when he reports risk aversion σ' and gets a contract (β, t) is

$$U_i(\beta, t, \sigma') = E\left[u_{\sigma_i}(\beta Q + t)\right]. \tag{A18}$$

Consider a twice-continuous, positive, and decreasing function $\beta(\sigma')$. There exists a transfer function $t(\sigma')$ such that the menu of wage contracts $(\beta(\sigma'), t(\sigma'))$ implements a mechanism in which the managers truthfully report their risk aversion to choose a wage contract if the following condition holds:

$$\left[\frac{\partial}{\partial \sigma_i}\left(\frac{\partial U_i(\beta(\sigma'), t(\sigma'), \sigma')/\partial \beta}{\partial U_i(\beta(\sigma'), t(\sigma'), \sigma')/\partial t}\right)\right]_{\sigma'=\sigma_i} \frac{d\beta(\sigma_i)}{d\sigma_i} \geq 0. \tag{A19}$$

Here the second term is negative by assumption. From (A18), we get

$$\frac{\partial U_i(\beta(\sigma'), t(\sigma'), \sigma')/\partial \beta}{\partial U_i(\beta(\sigma'), t(\sigma'), \sigma')/\partial t} = \frac{E\left[Q u'_{\sigma_i}(\beta Q + t)\right]}{E\left[u'_{\sigma_i}(\beta Q + t)\right]} = \frac{1}{\beta}\left(\frac{E\left[w u'_{\sigma_i}(w)\right]}{E\left[u'_{\sigma_i}(w)\right]} - t\right)_{w=\beta Q + t},$$

so (A19) will hold if we can show that

$$\frac{\partial}{\partial \sigma_i}\frac{E\left[w u'_{\sigma_i}(w)\right]}{E\left[u'_{\sigma_i}(w)\right]} \leq 0.$$

Upon expanding, we get

$$\begin{aligned}&\frac{\partial}{\partial \sigma_i}\frac{E\left[w u'_{\sigma_i}(w)\right]}{E\left[u'_{\sigma_i}(w)\right]}\\ &= \frac{1}{\left(E\left[u'_{\sigma_i}(w)\right]\right)^2}\left\{E\left[w\frac{d^2 u_{\sigma_i}}{d\sigma_i\, dw}\right]E\left[\frac{du_{\sigma_i}}{dw}\right] - E\left[w\frac{du_{\sigma_i}}{dw}\right]E\left[\frac{d^2 u_{\sigma_i}}{d\sigma_i\, dw}\right]\right\},\end{aligned}$$

which is negative by Chebyshev's inequality (see Mitrinovic and Vaic (1970), Theorem 10, p. 40) if

$$\frac{d^2 u_{\sigma_i}/d\sigma_i\, dw}{du_{\sigma_i}/dw} \tag{A20}$$

is decreasing in w. Differentiating (A20) with respect to w, the sufficient condition for this is

$$\frac{d^3 u_{\sigma_i}}{d\sigma_i\, dw^2}\frac{du_{\sigma_i}}{dw} - \frac{d^2 u_{\sigma_i}}{d\sigma_i\, dw}\frac{d^2 u_{\sigma_i}}{dw^2} < 0,$$

which can be obtained by expanding (14). Q.E.D.

Proof of Lemma 3: Let p^* be the threshold portfolio quality that the optimal contract implements for accepting a portfolio. The CEO accepts all portfolios with assessed quality of at least p^*. The portfolio accepted at the threshold provides at least as much expected payoff to the shareholders as the value from portfolio rejection (see Proposition 4) and the shareholders' expected payoff for accepted portfolios increases linearly with portfolio quality since their expected payoff from a good portfolio, $\Sigma x_i + h - W_H$, exceeds that from a bad portfolio, $\Sigma x_i + l - W_L$. Thus, the expected payoff to shareholders is a convex function of assessed portfolio quality. A signal s with precision q results in assessed portfolio quality $sq + 0.5(1 - q)$. While greater precision does not change the expected value of the assessed portfolio quality, it increases the variability (informativeness) of the assessed portfolio quality. By Jensen's theorem, this increases the shareholders' expected payoff, which is the expectation of a convex function of assessed portfolio quality. If $\Pi(q)$ is shareholders' expected payoff from the portfolio when the signal precision is q, then $\Pi'(q) > 0$.

For the second part of the lemma, suppose the true precision of the CEO's original signal is q_1. A rational CEO will invest in information production to achieve precision q_2 such that

$$\Pi'(q_2) - \theta'(q_2) = 0. \tag{A21}$$

An overconfident CEO believes the precision of her signal is $q_1(C) > q_1$ but has the same first-order condition, (A21), for the choice of precision. She invests $\theta(q_2) - \theta(q_1(C))$, which is less than the amount $\theta(q_2) - \theta(q_1)$ that a rational CEO would have invested, and she thus ends up with a signal with precision less than q_2. Q.E.D.

Proof of Proposition 8: An overconfident CEO underinvests less than a rational CEO and increases firm value. Since the rational CEO underinvests, the marginal benefit of overconfidence, calculated as $V'(C) > 0$ for $1 < C < C^*$ in the proof of Proposition 5, is strictly positive. An overconfident CEO underinvests in information production compared to a rational CEO and this reduces firm value. From the proof of Lemma 3, a rational CEO with $C = 1$ invests the optimal amount (see (A21)), so the marginal cost of underinvestment in information production is zero at $C = 1$. Hence, the marginal benefit of overconfidence exceeds the marginal cost at $C = 1$. Q.E.D.

Proof of Lemma 4:

(a) Constraint (18) can be eliminated as it follows from substituting (17) in (19) with $s = s^*(\hat{C})$. Constraint (19) is most restrictive for the highest signal $s = 1$, because with a very high signal the CEO has strong aversion to rejecting the portfolio and being considered excessively diffident. Thus, the board's problem reduces to

$$\begin{aligned}
&\underset{s^*(C), W_H, W_L, W_R}{\text{Max}} E\left[\sum x_i\right] \\
&\quad + \Pr(s \geq s^*(C)) \times E\left[p(s,1)(h - W_H) + \{1 - p(s,1)\}(l - W_L) \,|\, s \geq s^*(C)\right] \\
&\quad + \Pr(s < s^*(C)) \times (r - W_R) + E[\max(Y_2(\psi,\mu), Y_2(\zeta,\nu))]
\end{aligned} \tag{A22}$$

subject to

$$u_A(s^*(\hat{C}), \hat{C}) = u_R(s^*(\hat{C}), \hat{C}) \ \forall \hat{C}, \tag{A23}$$

$$u_R(1, \hat{C}) \geq u_L(1, \hat{C}) \ \forall \hat{C}, \tag{A24}$$

and

$$\begin{aligned} &\Pr(s \geq s^*(C)) E\left\{u_A(s, C) \,\middle|\, s \geq s^*(C)\right\} \\ &\quad + \Pr(s < s^*(C)) E\left\{u_R(s, C) \,\middle|\, s < s^*(C)\right\} \geq U_1. \end{aligned} \tag{A25}$$

Let $\eta(\hat{C})$, $\kappa(\hat{C}) \leq 0$, and $\delta \leq 0$ be the Lagrange multipliers for constraints (A23), (A24), and (A25), respectively, in the above problem. The IR constraint (A25) must bind because otherwise expected CEO wage can be reduced without violating the constraints by lowering $U(W_H)$, $U(W_L)$, and $U(W_R)$ by the same small constant. Next, to prove $\eta(\hat{C}) < 0$, consider the optimal solution to a simpler problem, obtained from the above problem by removing constraint (A23). If (A24) binds, the solution is

$$\begin{aligned} u(W_L) = u(W_R) &+ E\left\{U_2(\zeta, \nu) + B \times \Pr\left[Y_2(\zeta, \nu) > Y_2(\psi, \mu)\right] \middle| 1 < s^*(C)\right\} \\ &- E\left\{U_2(\zeta, \nu) + B \times \Pr\left[Y_2(\zeta, \nu) > Y_2(\psi, \mu)\right] \middle| 1 \geq s^*(C)\right\}, \end{aligned}$$

and

$$\begin{aligned} \frac{u'(W_H)}{u'(W_R)} &= \frac{E\left[p(s, 1) \mid s \geq s^*(C)\right]}{E\left[p(s, C) \mid s \geq s^*(C)\right]} \\ &\quad \times \frac{\Pr(s \geq s^*(C)) E\left[p(s, 1) \mid s \geq s^*(C)\right] + \Pr(s < s^*(C))}{\Pr(s \geq s^*(C)) E\left[p(s, 1) \mid s \geq s^*(C)\right] \dfrac{u'(W_R)}{u'(W_L)} + \Pr(s < s^*(C))}. \end{aligned}$$

If (A24) does not bind, the solution to (A22) and (A25) is

$$\frac{u'(W_H)}{u'(W_R)} = \frac{E\left[p(s, 1) \mid s \geq s^*(C)\right]}{E\left[p(s, C) \mid s \geq s^*(C)\right]}, \quad \frac{u'(W_L)}{u'(W_R)} = \frac{1 - E\left[p(s, 1) \mid s \geq s^*(C)\right]}{1 - E\left[p(s, C) \mid s \geq s^*(C)\right]}.$$

In both cases, as long as the CEO's expected deviation from rationality ($C = 1$) is sufficiently small, W_H and W_L are close enough so that

$$u_H(s, \hat{C}) - u_L(s, \hat{C}) < \chi, \tag{A26}$$

where χ satisfies

$$\chi \leq c / p(s^*(\hat{C}), \hat{C}) \ \forall \hat{C}. \tag{A27}$$

Multiplying (A26) by $p(s^*(\hat{C}), \hat{C})$, subtracting from (19) with $s = s^*(C)$, and substituting (24) results in

$$u_A(s^*(\hat{C}), \hat{C}) < u_R(s^*(\hat{C}), \hat{C}) - c + \chi p(s^*(\hat{C}), \hat{C}) < u_R(s^*(\hat{C}), \hat{C}),$$

where the last inequality follows from (A27). Thus, the right-hand side of (A23) exceeds the left-hand side and we must have $\eta(\hat{C}) < 0$ in the solution of (A22) to (A25). Let Γ be the expected payoff of the shareholders. The first-order condition for $s^*(\hat{C})$ to be the equilibrium threshold signal value for a CEO with overconfidence $\hat{C}$ is

$$\frac{\partial \Gamma}{\partial s^*(\hat{C})} - \eta(\hat{C})(u(W_H) - u(W_L))\frac{\partial p(s^*(\hat{C}), \hat{C})}{\partial s^*(\hat{C})}$$
$$- \delta\left\{u_A(s^*(\hat{C}), \hat{C}) - u_R(s^*(\hat{C}), \hat{C})\right\}\frac{d}{ds^*(\hat{C})}\Pr[s \le s^*(\hat{C})] = 0.$$

Substituting $\eta(\hat{C}) < 0$, (18), and (A23), we get

$$\frac{\partial \Gamma}{\partial s^*(\hat{C})} < 0.$$

Thus, the shareholders' expected payoff will increase if the CEO decreases the threshold signal value for portfolio acceptance. Finally, to verify that the beliefs in (21) to (24) are rational, we must show that $s^*(C)$ is a decreasing function. This follows from (18) and the fact that the CEO's second-period utility U_2 is single-peaked in perceived overconfidence, so rejecting a portfolio with a higher signal s causes the CEO to be considered more diffident and imposes a higher penalty than rejecting a portfolio with a lower s. If a CEO with low overconfidence is indifferent between accepting and rejecting a portfolio, then a CEO with higher overconfidence who estimates a higher probability of success will strictly prefer accepting the portfolio to rejecting it and hence must have a lower threshold signal for portfolio acceptance.

(b) The proof is identical to that of Lemma 2.

(c) Follows from part (a) and the fact that the threshold signal value for portfolio acceptance, $s^*(C)$, is decreasing in CEO overconfidence C. Q.E.D.

Proof of Lemma 5: The posterior distribution $\zeta(A_0 \mid s, \omega, X)$ of the CEO's ability when the board observes signal s, the portfolio outcome ω, and the project-specific payoffs X is

$$\zeta(A_0 \mid s, \omega, X) = \frac{\psi(A_0)\prod \xi(x_i, A_0)}{\displaystyle\int \prod \xi(x_i, A)\psi(A)\, dA}.$$

The posterior likelihood ratio of the probability that the CEO has ability A_H to the probability that the CEO has ability $A_L < A_H$,

$$\frac{\zeta(A_H \mid s, \omega, X)}{\zeta(A_L \mid s, \omega, X)} = \frac{\psi(A_H)\prod \xi(x_i, A_H)}{\psi(A_L)\prod \xi(x_i, A_L)},$$

is independent of s and ω and is increasing in x_i because of the monotone likelihood ratio property of $\xi(x_i, A_0)$. We now examine the board's posterior distribution $\nu(C \mid s, \omega, X)$ over the CEO's overconfidence. If the CEO rejects the portfolio, the posterior beliefs are given by

$$\nu(C \mid s, R, X) = \frac{1_{s<s^*(C)}\mu(C)}{\displaystyle\int_{C'} 1_{s<s^*(C')}\mu(C')\,dC'} = \frac{1_{C<s^{*-1}(s)}}{\displaystyle\int_{C'<s^{*-1}(s)} \mu(C')\,dC'}\mu(C).$$

If the CEO accepts the portfolio, the posterior beliefs are given by

$$\nu(C \mid s, H, X) = \nu(C \mid s, L, X) = \frac{1_{s\geq s^*(C)}\mu(C)}{\displaystyle\int_{C'} 1_{s\geq s^*(C')}\mu(C')\,dC'} = \frac{1_{C\geq s^{*-1}(s)}}{\displaystyle\int_{C'\geq s^{*-1}(s)} \mu(C')\,dC'}\mu(C).$$

When the portfolio is rejected, the posterior distribution of CEO overconfidence with a lower signal s_1 first-order-stochastically dominates the distribution with a higher signal $s_2 > s_1$ because the two distributions satisfy the monotone likelihood ratio property, given that the ratio

$$\frac{\nu(C \mid s_1, R, X)}{\nu(C \mid s_2, R, X)} = \frac{1_{C<s^{*-1}(s_1)}}{1_{C<s^{*-1}(s_2)}} \frac{\displaystyle\int_{C'<s^{*-1}(s_2)} \mu(C')\,dC'}{\displaystyle\int_{C'<s^{*-1}(s_1)} \mu(C')\,dC'}$$

is nondecreasing in C. When the portfolio is accepted, the posterior distribution of CEO overconfidence with a lower signal s_1 first-order-stochastically dominates the distribution with a higher signal $s_2 > s_1$ because the two distributions satisfy the monotone likelihood ratio property, given that the ratio

$$\frac{\nu(C \mid s_1, H, X)}{\nu(C \mid s_2, H, X)} = \frac{\nu(C \mid s_1, L, X)}{\nu(C \mid s_2, L, X)} = \frac{1_{C\geq s^{*-1}(s_1)}}{1_{C\geq s^{*-1}(s_2)}} \frac{\displaystyle\int_{C'\geq s^{*-1}(s_2)} \mu(C')\,dC'}{\displaystyle\int_{C'\geq s^{*-1}(s_1)} \mu(C')\,dC'}$$

is nondecreasing in C. The posterior distribution of CEO overconfidence when the portfolio is accepted first-order-stochastically dominates the prior distribution because the two distributions satisfy the monotone likelihood ratio property, given that the ratio

$$\frac{\nu(C \mid s, H, X)}{\mu(C)} = \frac{\nu(C \mid s, L, X)}{\mu(C)} = \frac{1_{C\geq s^{*-1}(s)}}{\displaystyle\int_{C'\geq s^{*-1}(s)} \mu(C')\,dC'}$$

is nondecreasing in C. The posterior distribution of CEO overconfidence when the portfolio is rejected is first-order-stochastically dominated by the prior distribution because the two distributions satisfy the monotone likelihood ratio property, given that the ratio

$$\frac{\mu(C)}{\nu(C \mid s, R, X)} = \frac{\displaystyle\int_{C' < s^{*-1}(s)} \mu(C')\, dC'}{1_{C < s^{*-1}(s)}}$$

is nondecreasing in C. Q.E.D.

Proof of Proposition 9: The CEO is retained if $Y_2(\zeta, \nu) \geq Y_2(\psi, \mu)$ and replaced otherwise, where Y_2 is increasing in the first argument and single-peaked in the second argument when the arguments are point distributions. Thus, the board's decision depends on its beliefs about the CEO's overconfidence and ability. First, consider a fixed signal value and portfolio outcome. Lemma 5 shows that these completely determine the board's beliefs ν about the CEO's overconfidence, so the CEO will be retained if the CEO's ability is inferred to be sufficiently high, and Lemma 5 shows that this happens if the project-specific payoffs X are sufficiently high. Next, consider a fixed X. Lemma 5 shows that this completely determines the board's beliefs about the CEO's ability, so the CEO will be replaced if the CEO's degree of confidence is too high or too low. From Lemma 5, when the portfolio is rejected, the CEO is likely to be less confident than a replacement CEO, so the incumbent CEO will be fired if she is evaluated to be sufficiently diffident. From Lemma 5, this occurs when the signal s is sufficiently high. When the portfolio is accepted, Lemma 5 shows that the CEO is likely to be more confident than a replacement CEO, so the incumbent CEO will be fired if she is evaluated to be sufficiently overconfident. From Lemma 5, this occurs when the signal s is sufficiently low. Q.E.D.

REFERENCES

Adams, Renée, and Daniel Ferreira, 2007, A theory of friendly boards, *Journal of Finance* 62, 217–250.

Almazan, Andres, and Javier Suarez, 2003, Entrenchment and severance pay in optimal governance structures, *Journal of Finance* 58, 519–547.

Amihud, Yakov, and Baruch Lev, 1981, Risk reduction as a managerial motive for conglomerate mergers, *Bell Journal of Economics* 12, 605–617.

Barber, Brad, and Terrance Odean, 2001, Boys will be boys: Gender, overconfidence and common stock investment, *Quarterly Journal of Economics* 116, 261–292.

Beetsma, Roel M.W.J., and Peter C. Schotman, 2001, Measuring risk attitudes in a natural experiment: Data from the television game show Lingo, *Economic Journal* 111, 821–848.

Ben-David, Itzhak, John R. Graham, and Campbell R. Harvey, 2007, Managerial overconfidence and corporate policies, Working paper, Duke University and Ohio State University.

Bernardo, Antonio E., and Ivo Welch, 2001, On the evolution of overconfidence and entrepreneurs, *Journal of Economics & Management Strategy* 10, 301–330.

Bhattacharya, Sudipto, and Dilip Mookherjee, 1986, Portfolio choice in research and development, *RAND Journal of Economics* 17, 594–605.

Coval, Joshua, and Anjan V. Thakor, 2005, Financial intermediation as a beliefs-bridge between optimists and pessimists, *Journal of Financial Economics* 75, 533–570.

Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam, 1998, Investor psychology and security market under- and over-reactions, *Journal of Finance* 53, 1839–1885.

Dasgupta, Partha, and Eric Maskin, 1987, Simple economics of research portfolios, *Economic Journal* 97, 581–595.

de Meza, David, and Clive Southey, 1996, The borrower's curse: Optimism, finance and entrepreneurship, *Economic Journal* 106, 375–386.

DeBondt, W. F. M., and R. H. Thaler, 1995, Financial decision-making in markets and firms: A behavioral perspective, in R. Jarrow, V. Maksimovic, and W. Ziemba, eds.: *Finance, Handbook in Operations Research and Management Science, Vol. 9* (Amsterdam: Elsevier).

Diamond, Peter A., and Joseph E. Stiglitz, 1974, Increases in risk and risk aversion, *Journal of Economic Theory* 8, 337–360.

Fudenberg, Drew, and Jean Tirole, 1991, *Game Theory* (Cambridge, MA: The MIT Press).

Gervais, Simon, J. B. Heaton, and Terrance Odean, 2007, Overconfidence, investment policy, and manager welfare, Working paper, Duke University.

Grinblatt, Mark, and Matti Keloharju, 2008, Sensation seeking, overconfidence and trading activity, *Journal of Finance* (forthcoming).

Grossman, Sanford J., and Oliver D. Hart, 1983, An analysis of the principal-agent problem, *Econometrica* 51, 7–45.

Grund, Christian, and Dirk Sliwka, 2006, Performance pay and risk aversion, IZA Discussion paper No. 2012, University of Bonn.

Guiso, Luigi, and Tullio Jappelli, 2006, Information acquisition and portfolio performance, CEPR Discussion paper 5901.

Harris, Milton, and Artur Raviv, 2008, A theory of board control and size, *Review of Financial Studies* 21, 1797–1832.

Hermalin, Benjamin E., 1998, Towards an economic theory of leadership: Leading by example, *American Economic Review* 88, 1188–1206.

Hermalin, Benjamin E., 2005, Trends in corporate governance, *Journal of Finance*, 60, 2351–2384.

Hermalin, Benjamin E., and Michael Weisbach, 1998, Endogenously chosen boards of directors and their monitoring of the CEO, *American Economic Review* 88, 96–118.

Hermalin, Benjamin E., and Michael Weisbach, 2003, Boards of directors as an endogenously determined institution: A survey of the economic literature, *Economic Policy Review* 9, 7–26.

Hirshleifer, David, 2001, Investor psychology and asset pricing, *Journal of Finance* 56, 1533–1597.

Holmstrom, Bengt, 1999, Managerial incentive problems: A dynamic perspective, *Review of Economic Studies* 66, 169–182.

Holmstrom, Bengt, and Joan Ricart i Costa, 1986, Managerial incentives and capital management, *Quarterly Journal of Economics* 101, 835–860.

Jewitt, Ian, 1989, Choosing between risky prospects: The characterization of comparative statics results, and location independent risk, *Management Science* 35, 60–70.

Kahneman, Daniel, Paul Slovic, and Amos Tversky, eds., 1982, *Judgment under Uncertainty: Heuristics and Biases* (Cambridge University Press).

Khambata, Dara, and Wei Liu, 2006, Cultural dimensions, risk aversion and corporate dividend policy, *Journal of Asia-Pacific Business* 6, 31–43.

Klette, Tor, and David de Meza, 1986, Is the market biased against risky R&D? *RAND Journal of Economics* 17, 133–139.

Kyle, Albert S., and Albert F. Wang, 1997, Speculation duopoly with agreement to disagree: Can overconfidence survive the market test? *Journal of Finance* 52, 2073–2090.

Larwood, Laurie, and William Whittaker, 1977, Managerial myopia: Self-serving biases in organizational planning, *Journal of Applied Psychology* 62, 194–198.

Lazear, Edward P., and Sherwin Rosen, 1981, Rank-order tournaments as optimum labor contracts, *Journal of Political Economy* 89, 841–864.

Lorsch, Jay W., and Elizabeth MacIver, 1989, *Pawns or Potentates: The Reality of America's Corporate Boards* (Harvard Business School Press).

Malmendier, Ulrike, and Geoffrey Tate, 2005, CEO overconfidence and corporate investment, *Journal of Finance* 60, 2661–2700.

Malmendier, Ulrike, and Geoffrey Tate, 2008, Who makes acquisitions? CEO overconfidence and the market's reaction, *Journal of Financial Economics* (forthcoming).

Milbourn, Todd, Richard Shockley, and Anjan V. Thakor, 2001, Managerial career concerns and investments in information, *RAND Journal of Economics* 32, 334–351.

Milgrom, Paul R., 1981, Good news and bad news: Representation theorems and applications, *RAND Journal of Economics* 12, 380–391.

Mitrinovic, D. S., and P. M. Vaic, 1970, *Analytic Inequalities* (New York: Springer Verlag).

Pattillo, Catherine, and Mans Söderbom, 2000, Managerial risk attitudes and firm performance in Ghanaian manufacturing: An empirical analysis based on experimental data, CSAE Working paper 2000–17, University of Oxford.

Prendergast, Canice, 1999, The provision of incentives in firms, *Journal of Economic Literature* 7, 7–63.

Ramakrishnan, Ram T. S., and Anjan V. Thakor, 1984, Information reliability and a theory of financial intermediation, *Review of Economic Studies* 51, 344–351.

Ramakrishnan, Ram T. S., and Anjan V. Thakor, 1991, Cooperation versus competition in agency, *Journal of Law, Economics and Organization* 7, 248–283.

Rogerson, William P., 1985, The first-order approach to principal-agent problems, *Econometrica* 53, 1357–1368.

Rothschild, Michael, and Joseph Stiglitz, 1976, Equilibrium in competitive insurance markets: An essay on the economics of imperfect information, *Quarterly Journal of Economics* 90, 629–649.

Sandroni, Alvaro, and Francesco Squintani, 2007, Overconfidence, insurance, and paternalism, *American Economic Review* 97, 1994–2004.

Shleifer, Andrei, and Robert Vishny, 1997, A survey of corporate governance, *Journal of Finance* 52, 737–783.

Song, Fenghua, and Anjan V. Thakor, 2006, Information control, career concerns, and corporate governance, *Journal of Finance* 61, 1845–1896.

Suto, Megumi, Lukas Menkhoff, and Daniela Beckmann, 2005, Behavioral biases of institutional investors under pressure from customers: Japan and Germany vs. U.S., Waseda University Institute of Finance Working paper WIF-05-006, Japan.

Van Den Steen, Eric, 2004, Rational optimism (and other biases), *American Economic Review* 94, 1141–1151.

Van Den Steen, Eric, 2005, Organizational beliefs and managerial vision, *Journal of Law, Economics, and Organization* 21, 256–283.

Wang, Albert, 2001, Overconfidence, investor sentiment, and evolution, *Journal of Financial Intermediation* 10, 138–170.

Weinstein, Neil D., 1980, Unrealistic optimism about future life events, *Journal of Personality and Social Psychology* 39, 806–820.

Wu, Brian, and Ann Marie Knott, 2006, Entrepreneurial risk and market entry, *Management Science* 52, 1315–1330.