Advice, Information and the Reputation of CEOs*

Nina Walton
UCLA

November 5, 2007

Abstract

I investigate when and how principals can use external experts to monitor agents when agents control information necessary for monitoring. In particular, I provide a theoretical model to examine in what circumstances external experts can assist boards of directors in discharging their management oversight and investment approval duties. Management controls relevant information about an investment but may not have the ability to process it to make an informed recommendation. External experts have the ability, but may not enjoy access to this information. I demonstrate how using experts to certify management recommendations affects the disclosure incentives of management. Because certification serves as a signaling mechanism, when managers have the incentive to truthfully reveal all that they know, certification mandates are unnecessary since managers will choose to seek out second opinions. When information disclosures cannot be easily verified, certification mandates can be counterproductive, elevating the status of costly second opinions that always agree with management recommendations. In the absence of incentives for truthful disclosure, it is better for boards to allow management to determine when to seek second opinions, or alternatively, if second opinions are required to allow management and advisors to pool their recommendations.

*I am grateful to David Levine, Jean-Laurent Rosenthal and Bill Zame for their ongoing support, advice and patience. In addition, I thank Hugo Hopenhayn, Jernej Copic, Antonio Bernardi, Mark Garmaise, Anjan Thakor and participants in seminars at UCLA and Washington University in St Louis for their insightful comments; and Tracy Francis, Neal Lerner, and Tamir Maltz for helpful conversations about their own professional practices as advisors to firms. All mistakes are, of course, my own.

1 Introduction

When and how can principals use external experts to monitor agents, when agents control access to information necessary to monitor? What trade-offs do principals face between oversight of agents and information disclosure by agents? While the monitoring capacity of principals has relevance in many economic settings, of recent concern is the efficacy of board oversight of management when management controls private information about the firm it runs. With the passage of the Sarbanes-Oxley Act of 2002, Congress instituted a set of new rules governing the extent of board independence, the role of independent experts and the quality of management disclosures. Sarbanes-Oxley aims to improve outside monitoring of management, with the goal of improving firm performance.

In this paper, I make a distinction between *oversight*, where principals actively evaluate agents, and *disclosure*, where agents reveal private information. I argue that improved oversight may not lead to improved disclosure for two main reasons. Firstly, provisions that mandate increased board oversight ignore the incentives for managers to behave properly. Good managers do not need oversight but would like to be evaluated, and so are willing to comply with rules and disclose information. Bad managers, who do need oversight, would rather not reveal relevant information since they may be penalized for doing so. Secondly, provisions strengthening the role and independence of outside auditors and other expert advisors who assist boards by certifying management recommendations, fail to appreciate that the quality of outside certifications is only as good as the quality of the information they are based on and that management can often manipulate this information. Put together, these reasons imply that provisions intensifying board oversight through an external certification mechanism will only be successful if management has enough of an incentive to truthfully disclose all that it knows.

I demonstrate that when such incentives exist, certification mandates are unnecessary since managers will choose to institute certification mechanisms on their own. In the absence of disclosure incentives, certification mandates may be counterproductive. To evaluate the benefits of regulation augmenting board oversight and management disclosure obligations (as contained in Sarbanes-Oxley), I build a theoretical model in which there is a CEO who is concerned about his reputation for competence with the board, and a third party advisor who is known to be both competent and honest. The CEO must recommend to the board whether to make an investment or not. The CEO controls access to internal information relevant to the investment, but he may or may not have the ability to adequately process it to make an informed recommendation. On the other hand, the advisor is known to have the ability, but must rely upon the CEO for access.

I show that within the context of this model, to the extent that management reveals all relevant information, obtaining outside advice can be an effective

means of improving decision making within corporations. However, and as is likely, when CEOs can filter information strategically, a second opinion will give little or no information to the board about the appropriateness of a CEO's decision. In fact, relying on such opinions may be detrimental, since advisors will receive information that leads their opinions to always validate the proposed course of action. Second opinions based on inaccurate information therefore give undue credence to management claims, and mandating certifications in these circumstances only exacerbates principal-agent issues. When information quality is likely to be low, the board can partially restore the value of second opinions by reducing its direct oversight of the CEO, allowing the CEO to control the relationship with the advisor directly. By effectively renouncing its intent to use the recommendation to evaluate the CEO, the board reduces the CEO's incentive to disclose information strategically in an effort to improve his own reputation. Alternatively, the board can invest in a technology which increases the likelihood that CEO disclosures are accurate.

In the absence of rules requiring certification, CEOs who want to impress their boards (who in turn, determine CEO compensation) will internalize the weight boards give to the certification in deciding whether or not to obtain a second opinion. When information disclosures are reliable, boards should give weight to second opinions. Competent CEOs will then use advisors as a matter of course, since certification serves as a powerful signaling mechanism to the board. On the other hand, if information cannot be reliably transmitted to experts, the board should place no weight on the certification, and CEOs can then more effectively signal their competence by *not* using an advisor. Therefore, measures that increase the reliability of information coming from management also ensure that policies supporting certification achieve their desired ends, while reducing the need for such rules in the first place. In summary, credible disclosures facilitate oversight, but oversight does not necessarily facilitate credible disclosures.

In general, principal-agent models do not make a distinction between oversight of agents and agent disclosures of private information. Instead it is usually assumed that principals *monitor* agents, which involves instituting both oversight and disclosure mechanisms. Distinguishing between oversight and disclosure is useful however, when the information disclosed will have an effect not only on the agent's reputation, but also on firm profitability. When private information has this dual purpose, principals may prefer to sacrifice oversight (or evaluation) capacity if it means higher quality disclosures. Further, hierarchical agency models, with a principal, agent and supervisor, tend to assume that at a cost, supervisors (or auditors) can access all relevant information. In reality, agents are often better informed than both principals and supervisors, and moreover, control access to information relevant to their own performance. While some information will no doubt be verifiable, there may be other relevant information which is "soft" and unverifiable.

There are many economic settings where information largely controlled by agents is relevant for both agent reputation and principal performance. For example, loan officers or brokers rely mainly on the information applicants themselves provide in determining whether to grant a loan. Insurance brokers depend on the honesty of potential clients in making a coverage decision. Recruitment agencies depend on job seekers to supply pertinent information about career background and skills. While some information will be verifiable, other information will be within the sole control of the applicant, who decides how much of it to reveal knowing that doing so will affect their chances at getting a loan, insurance or a job respectively. In determining whether to reveal all information, agents will take into account how the principal intends to use the information, the likelihood that the supervisor will detect that the agent is not fully revealing all that he knows, and the penalties the agent will suffer upon detection.

Likewise, boards of directors of corporations are tasked with making CEO hiring, retention and compensation decisions based on their perception of CEO ability and firm performance. The Sarbanes-Oxley Act (hereafter SOX) tried to address informational problems that arise in publicly traded corporations, although the controversy over how best to monitor management, as well as potential remedies, is almost as old as the corporate form itself.[1] SOX introduced new mandates that increased the independence of boards and its advisors from management[2], and enhanced the obligations of managers to certify the veracity of disclosures of internal information to outsiders.[3] Hailed as "the most far-reaching reforms of American business practices since the time of Franklin D. Roosevelt"[4], SOX legislation is controversial because it creates a new system of

[1] Debates over how to induce management to reveal information are universal and long-standing. For example, in the United Kingdom, to overcome disclosure problems, the Joint Stock Companies Act (1844) and the Companies Clause Consolidation Act (1945) required annual statutory audits of corporations, with the auditor being elected by shareholders and not employed by the company. In France, the legislature passed a law in 1857 with the explicit intent of preventing fraud and misrepresentation by firm insiders. The committee head, in his report to the legislature, stated that the bill "attempts as far as possible, to replace deceit and lies with truth and fairness". The legislature considered and rejected mandating external certification of internal claims, relying on *councils of surveillance* (analogous to boards of directors) instead, spelling out their duties and making them liable for failure (Freedeman (1979)). Legislative efforts to ensure quality disclosures in the United States occurred much later. Outside verification of internal claims was only mandated for newly incorporated companies by the NYSE in the 1920's, however, it was not until the passage of the Securities Acts of 1933 and 1934 that external auditing was required for all publicly traded corporations. Until then, disputes over disclosures were mainly dealt with by state courts. See Coffee (2006).

[2] For example, section 301 requires all members of the audit committee to be independent.

[3] Sections 302, 303, 401, 404 and 409 relate to the adequacy of management disclosures to investors. Section 302 requires that senior corporate officers individually certify that their financial statements accurately and fairly present the company's financial position and operating results. Section 303 provides for an SEC rule prohibiting fraudulently influencing or misleading an auditor for purposes of rendering financial statements misleading. Section 404 also makes management responsible for reporting on the state of the company's internal control systems. Section 409 requires management to make timely disclosures of material changes in the company's financial condition.

[4] President George W. Bush, as reported by Elisabeth Bumiller: "Bush Signs Bill Aimed

federal regulatory mandates that override alternative contractual arrangements between managers, boards and shareholders (Ribstein (2002), Romano (2006), Coates (2007)). Supporters argue that mandates are necessary to restore confidence in, and integrity of, the public financial markets. Detractors argue that costs of SOX's implementation, in particular those associated with assuring the quality of disclosures about internal information, outweigh the corresponding benefits.[5] The current debate over SOX and its many predecessors, raises an important puzzle: when and how can increased oversight and disclosure requirements improve corporate returns?

Evidence on the relationship between board independence (usually equated with oversight) and firm performance is mixed: for the most part independent boards do not seem to result in better firm performance than firms dominated by insiders.[6] One potential reason for the lack of a clear-cut benefit is that independent boards suffer from informational deficiencies. In addition to being time and resource constrained, independent boards, by definition, are not deeply involved with the internal affairs of the corporations they oversee. While independence may lower the risk of conflicts of interest, it also reduces the knowledge board members have about the firm's internal affairs.

Because of this lack of knowledge, in order to satisfy its fiduciary obligations to act diligently in approving investments and monitoring management, boards very often rely on third-party expert advisors. Third party advisors include auditors, accountants, attorneys, consultants, and investment banks. Investors also rely on ratings agencies and stock analysts, who advise on the value of firm investment strategies. As "gatekeepers" or "reputational intermediaries", independent advisors can be credible purveyors of impartial advice, because in advising their clients, they place their own reputation for expertise on the line. Theory suggests that longer-term concerns about one's reputation for expertise should be sufficient to overcome any short-term benefits from knowingly giving incorrect or bad advice (Kreps and Wilson (1982), Milgrom and Roberts (1982), Fudenberg and Levine (1989)). By certifying management disclosures and recommendations, advisors can thus credibly communicate that management is acting in the best interests of the firm (Coffee (2006), Kraakman (1985)).[7]

---

at Fraud in Corporations", The New York Times, July 31, 2002, page A1.

[5] Section 404 is the most controversial provision because it has proven to be the most costly. Section 404 requires management and outside auditors to report on the scope and adequacy of the company's internal control system. According to *COSO's Integrated Framework*, internal control is broadly defined as a process, effected by an entity's board of directors, management and other personnel, designed to provide reasonable assurance regarding the achievement of objectives in the following categories: 1) Effectiveness and efficiency of operations, 2) Reliability of financial reporting, 3) Compliance with applicable laws and regulations. See http://www.coso.org/publications/executive_summary_integrated_framework.htm.

[6] There is some evidence that independent compensation, audit and nominating committees are associated with higher firm value (Chhaochharia and Grinstein (2004)). For theoretical models describing the endogeneity of board selection and performance, see Fama and Jensen (1983), Hermalin and Weisbach (1998), Almazan and Suarez (2003) and Warther (1998).

[7] While in the model below I assume that advisors act honestly and competently, there

Because independent advisors are outsiders to the firm, they cannot immediately access all relevant internal information. Instead, they must rely on management to give them the information they need to reach an informed and accurate opinion about a given proposal. The board can evaluate management by comparing its recommendation about a given investment with that of the advisor. Certification (when the advisor agrees with the manager's recommendation) facilitates oversight if the board expects the manager to be able to make an equally informed recommendation as the advisor (who is known to be an expert). When the board does not expect the manager to be able to reach the same conclusion as the advisor, no such comparison takes place. Therefore the role the advisor plays in assisting the board – certifier or specialist – depends on board expectations as well as the actual ability of the manager. When management's payoff depends on the esteem in which it is held by the board, the fact that the board can and will compare advisor and management recommendations changes management's incentives to disclose all relevant information to advisors.

This paper extends formal models showing how career concerns impact investment recommendations (Holmström (1999), Holmström and Ricart i Costa (1986)), by introducing the possibility of third party certification. Auditing has been considered in the literature on optimal contracting between principals and agents where principals are concerned with the effort level of their agents (Baron and Besanko (1984), Kofman and Lawarrée (1993), Khalil (1997)). Kofman and Lawarrée make a distinction between internal and external auditors, noting that internal auditors are costless, well informed but are likely to collude with agents, while external auditors are costly, not informed but will not engage in collusion. Baron and Besanko (1984) and Khalil (1997) do not have mandated audits, but instead describe the optimal audit policy when agents have private information about costs, or production levels respectively. Similar to these models, agents have private information, and auditing (or certification) is one mechanism that can be used by the principal to uncover it. By contrast, these models all focus on the incentives to audit agent activities *ex-post*, while I examine certification for management recommendations *ex-ante*, where the certification itself can be distorted by the private information transmitted by the agent.

Like others who have studied information distortion within organizational hierarchies (for example, Aghion and Tirole (1997), Levitt and Snyder (1997), Song and Thakor (2006), Adams and Ferreira (2006)), I have endeavored to provide a positive model describing information control within an organization,

is evidence that reputation alone may not be adequate to overcome short-term incentives to collude with management. Arthur Anderson's role in the Enron accounting scandal is a case in point. While questions about the incentives advisors have to collude with management, or to turn a blind eye to corporate malfeasance are interesting, they lie outside the scope of this paper which focuses on the incentives of management to disclose information and the effect of monitoring requirements on the value of outside advice. Policies that aim to reduce advisor incentives to collude with management belong to a different regulatory regime that those dealing with managerial incentives to withhold information from advisors.

and the agent's incentives to share that information with his principal, when the principal could later use it against him. Unlike these models, I incorporate the role that third parties outsiders play in assisting principals, examining specifically how regulatory solutions to problems of information distortion that utilize outside verification, may or may not be viable.

The remainder of the paper proceeds as follows. Section 2 presents the model. After first describing management's incentives to share information with advisors, Section 3 describes equilibrium results. Given the different disclosure regimes, the analysis compares outcomes where use of advisors is mandated with that where use is within managerial discretion. Section 4 concludes.

2 The Model

Consider an environment in which there is a single firm with many shareholders, a CEO who runs the company, an independent board of directors to whom the CEO reports, and an outside expert advisor who is available to provide advice to the company for a fee. To grow the firm's business, the CEO must make new investments, which must ultimately be approved by the board of directors. Although the board is independent, it relies on the CEO to make recommendations about investment opportunities, since it has neither the resources, nor the capability to come to a view on its own. In addition to approving large investments, the board evaluates CEO performance. This entails setting CEO compensation to reflect his ability at making investments, and replacing the CEO if the board loses confidence in his abilities.[8]

At the beginning of each period, an investment opportunity becomes available to the firm, of which both the board and the CEO are aware. If approved, the investment costs the firm $c \in (0, 1)$. A good investment (G) is profitable for the firm, yielding a return of 1, while a bad investment (B) has a return of 0, resulting in a net loss. Investment outcome is observed only if the investment is made. If the board does not approve the investment, it is not made, and the outcome is unknown. However, both the CEO and the board know that the unconditional likelihood that the investment outcome is good is $\theta \in (0, 1)$.

In making his recommendation whether to invest or not, the CEO has access to both private and public information about the company. Public information is available to everyone including the board. An example of public information is knowledge about the competitive landscape or the market in general. Access to

[8] The model focuses on CEO ability to make investments. Of course, CEOs carry out other tasks such as managing the firm's existing business. Many of these tasks however, require decisions about investments (such as investing in R&D to bring out a next-generation product) or divestments (whether to sell or shut down an existing business).

private information, on the other hand, is controlled by the CEO.[9] Importantly, controlling access means that information given to the board and other outsiders may only be partial, yet those outsiders would not be so aware without being explicitly informed. For example, in a proposed merger, the potential synergy from combining firm operations is private information. Because of the "soft" nature of this information, it may be difficult or impossible for outsiders to verify its accuracy or completeness prior to the merger's execution. While the board and the public may eventually discover exactly what the CEO knew at the time, at least in the short-run, they remain unaware of the full extent of the CEO's private information.

The CEO's *signal* consists of all pieces of private information that may be relevant to the investment. Each individual piece of information or element is indicative of the likely investment state. I assume that an element can either be "good" or "bad". A good element indicates that the investment is more likely to be good, while a bad element indicates that the investment is more likely to be bad. For this reason, we say that signals are either *ambiguous* (A) or *unambiguous* (U). In an unambiguous signal, the individual elements do not conflict with one another (that is, all elements are good, or all elements are bad), while in an ambiguous signal, there is conflict between the different elements (the set contains both good and bad elements).

Using his signal, the CEO draws an inference about the investment which is either good or bad. Based on this inference, the CEO makes a recommendation to the board about whether or not to make the investment. CEOs can have either *high ability* (H) or *low ability* (L) in drawing inferences about the true investment state. The relationship between CEO ability and information ambiguity can be described as follows. While both types always draw the same signal, they will only draw the same inference about an investment opportunity when the signal is unambiguous. For ambiguous signals, only the high type will be able to determine which element is relevant for the investment outcome. The low type cannot distinguish between the conflicting elements to determine which state is more likely. He draws an inference about the investment outcome that is completely noisy, giving no indication of the true investment state. Let s_i^θ denote the CEO's inference about the likely investment outcome, and $\widehat{s}_i^\theta$ denote the CEO's recommendation to the board, where $\theta \in \{G, B\}$ and $i \in \{H, L\}$.

While inferences may be informative, they will not always be correct. Since unpredictable events can sour even the most promising of investment opportunities, investment outcomes are inherently uncertain. Therefore, the most accurate inference will only be imperfectly correlated with the true investment state, and high types cannot be absolutely confident that they will be correct.

[9] The CEO relies on his management team who in practice control the information in question. To simplify the discussion, I refer to the CEO rather than top corporate officers, since the CEO either knows or could have known about internal information relevant to the investment.

The high type draws an accurate inference with probability p, regardless of the information set. To ensure that the high type's inference is more likely to be right than wrong, I assume that $p > \max\{\theta, 1-\theta\}$. When the signal is unambiguous, the low type also draws an accurate inference with probability p. For ambiguous signals, I assume without loss of generality, that the low type is equally likely to draw a good and bad inference regardless of the true state. Hence his signal is completely noisy and therefore uninformative. Summarizing,

$$\Pr\left(s^G \mid G, U\right) = \Pr\left(s^B \mid B, U\right) = p$$

$$\Pr\left(s_H^G \mid G, A\right) = \Pr\left(s_H^B \mid B, A\right) = p$$

$$\Pr\left(s_L^G \mid G, A\right) = \Pr\left(s_L^G \mid B, A\right) = \tfrac{1}{2}$$

The probability that the high type draws a good inference and the investment state is good is $p\theta$, and the probability he draws a good inference and the true investment state is bad is $(1-p)(1-\theta)$. Therefore, the unconditional probability he draws a good inference is $p\theta + (1-p)(1-\theta)$. Similarly, the probability the high type draws a bad inference when the true investment state is bad is $p(1-\theta)$, and the probability he draws a bad inference when the true state is good is $(1-p)\theta$. The unconditional probability the high type draws a bad inference is $p(1-\theta) + (1-p)\theta$. Using Bayes Rule, we can calculate the conditional probabilities that the high type is right and wrong given his inference:

$$\Pr\left(G \mid s_H^G\right) = \tfrac{p\theta}{p\theta+(1-p)(1-\theta)}$$

$$\Pr\left(B \mid s_H^G\right) = \tfrac{(1-p)(1-\theta)}{p\theta+(1-p)(1-\theta)}$$

$$\Pr\left(B \mid s_H^B\right) = \tfrac{p(1-\theta)}{p(1-\theta)+(1-p)\theta}$$

$$\Pr\left(G \mid s_H^B\right) = \tfrac{(1-p)\theta}{p(1-\theta)+(1-p)\theta}$$

The probabilities of good and bad signals are the same for low types who draw an unambiguous signals. Following an ambiguous signal, the probability the low type makes the correct recommendation is just the unconditional probability that the investment is good or bad. Therefore, we have:

$$\Pr\left(G \mid s_L^\theta\right) = \Pr\left(G\right) = \theta$$

$$\Pr\left(B \mid s_L^\theta\right) = \Pr\left(B\right) = 1-\theta$$

The board is uncertain about the CEO's type. It maintains a prior belief that the CEO has high ability with probability $\alpha \in (0,1)$. The board derives its belief from all that it knows about the CEO's performance in this and in other positions. The CEO however, is aware of his own ability type. One might alternatively think that there is an initial move by nature determining the probability distribution over CEO types in the population. The board knows these probabilities, but the actual type is known only to the CEO.

In addition to the CEO's recommendation, the board can rely upon a recommendation from a third party advisor, who charges the firm a fixed fee z. The advisor depends on the CEO to give her access to the signal from which she also draws an inference, denoted by s_a^θ. It is common knowledge that the advisor is honest, always recommending in accordance with her inference. She is also known to have high ability. I make the assumption that when presented with the same signal, high types all draw identical inferences. This assumption implies that uncertainty is systematic. High ability CEOs and advisors will therefore be incorrect about exactly the same factors when presented with exactly the same information.[10] Therefore, the advisor may be valuable to the board for two reasons: 1) she provides expert advice reassuring the board that an investment is worth making (or not), and 2) she increases the board's capacity to evaluate the CEO, by allowing the board to compare CEO and advisor recommendations. Going forward, we focus on CEO strategies following the receipt of an ambiguous signal only.[11] However, the possibility of drawing an unambiguous signal is important, because low types know how the advisor will recommend conditional upon receiving unambiguous information.

Certification as described here is somewhat analogous to investment models with reputational herding. Scharfstein and Stein (1990) provide a model where an investment manager (who does not know his own ability type but has access to private information) sees another manager's investment recommendation before making his own. Trueman (1994) also provided a model investigating how analysts (who do know their ability type) predict corporate earnings to protect their reputations.[12] Both of these models demonstrate that agents who are concerned about their reputations will ignore their own private information (even when they have high ability), preferring to recommend in a way that mimics

[10] The assumption of perfect correlation provides a simple structure in which to examine the demand for certification. The results below will follow through so long as inferences between advisors and high types have some degree of positive correlation. This is not the case when inferences are independent. Then, the advisor's recommendation does not give the board additional information about CEO type. The second opinion will, however, give the board additional information about the investment outcome, and therefore assists it in determining whether to approve the investment, even when the CEO has high ability.

[11] For unambiguous signals, there is no difference in ability types since both types draw the same informative inference. Since the game is one of signaling, we are interested in examining strategies in which high types are able to separate from low types. When there is essentially no difference between them, both types will behave identically in equilibrium.

[12] See Ottoviani and Sorensen (2000) for additional commentary on the relationship between correlation of private signals and private knowledge about one's type.

recommendations by others already made. While in our model, high types will not ignore their own private information about likely investment outcomes when making their recommendation, they will take into account the impact of an additional recommendation on their reputation. Just as in the herding models where a concurring prior recommendation protects analysts in the event investments are bad, so too does a concurring second opinion that is obtained after the CEO makes his own recommendation. In reputational herding models, the manager is not given a choice about the presence of additional recommendations. I show below, that if given a choice, under the assumptions outlined above, high types will do better with additional recommendations, than without. The second opinion certifying the CEO's recommendation to invest or not, gives more credibility to that recommendation, and provides the CEO with protection in the event that the investment performs poorly.[13]

Before play begins, at time $t = 0$, the board forms its belief about the likelihood that the CEO has high ability, and nature determines the true state of the investment. At time $t = 1$, the CEO uses his signal to draw an inference $s_i \in \left\{s^G, s^B\right\}$. After drawing the inference, a decision is made whether or not to hire an advisor. Let $a_i \in \left\{a^h, a^n\right\}$ denote the hiring decision where a^h represents the decision to hire, and a^n is the decision not to hire. The question of who makes this decision is left deliberately vague at this point, since we will compare outcomes under different scenarios (summarized in Figures 1a - 1c below) to evaluate the efficacy of enhanced board oversight.

The CEO discloses a signal to the advisor if she is hired. Since the CEO controls access to the signal, he determines which elements the advisor sees. I assume that the CEO cannot manufacture signals, but can strategically withhold elements that might be relevant to the recommendation. Specifically, the CEO has a choice to present the true ambiguous signal, or censor elements to present an unambiguous signal. The disclosure strategy, denoted by $\varphi_i\left(U \mid A\right)$, represents the probability that the CEO discloses an unambiguous signal conditional upon seeing an ambiguous signal. The choice is not innocuous. Recall that low types will only be guaranteed of drawing the same inference as high types and advisors when the signal is unambiguous.

If the CEO presents an unambiguous signal to the advisor when the true signal is actually ambiguous, there is always a chance he will be caught and held accountable for the misrepresentation. Accordingly, I assume that the advisor detects the misrepresentation with probability $0 \leq \rho \leq 1$, upon which discovery, she will inform the board. The detection probability ρ is designed to capture the amount of risk or liability that accrues to management from non-disclosure of material information.

[13] Of course, this result depends on the high type knowing that the additional recommendation is coming from another high ability type. Certification is only credible (and therefore desirable), when the certifying party's type is known. Below I investigate how low types deal with their ignorance about how the advisor will recommend, and how this ignorance increases the chances of certifications being unreliable.

Conditional on the advisor being hired, a choice must also be made about reporting structure: specifically, whether recommendations should be "pooled" or "separated". The notion of pooled and separated recommendations, not to be confused with pooling and separating strategies, is designed to capture the idea of certification – that is, whether the board can compare recommendations to evaluate the CEO. With pooled reports, the CEO sees the advisor's recommendation before making his own. Since high types and advisors will always draw the same inference (because of perfectly correlation), recommending differently to the advisor will lead the board to believe the CEO has low ability. Therefore, low types will never disagree with the advisor, and recommendations always agree. Boards can only compare CEO and advisor recommendations if the reports are separated in the sense that the CEO does not know how the advisor will recommend prior to reporting to the board. Separating recommendations by time – that is, hiring the advisor only *after* the CEO makes a recommendation – is one credible method of providing separated reports. Let $r \in \{r^s, r^p\}$ denote the reporting strategy where r^s represents separated reports while r^p represents pooled reports. If reports are separated or if the advisor is not hired, the CEO makes his recommendation known to the board at this time.

At time $t = 2$, the advisor (if hired) and the CEO (if reports are pooled) make their recommendations known to the board. The CEO uses a recommendation strategy denoted by $\sigma_i(\widehat{s} \mid s)$. This strategy is the probability the CEO makes a certain recommendation conditional on his inference. Given that the low type draws an uninformative inference when the signal is ambiguous, his recommendation strategy is independent of his inference, that is, $\sigma_L(\widehat{s} \mid s) = \sigma_L(\widehat{s})$. After seeing the recommendation(s), the board decides whether to approve the investment or not. The board is risk neutral and approves investments with the goal of maximizing expected profits for the firm, which is just the expected benefit of the investment, minus the cost of investing c, minus any advisory fees z that are incurred.

At time $t = 3$, the investment outcome is realized if the investment is made. At this point, the board updates its prior belief about CEO ability using Bayes rule. The posterior belief, denoted by $\widehat{\alpha}$, is equivalent to the CEO's compensation going into the next period.[14] Given the structure of the game, the CEO's actions in the first period, only affect his objective function through his continuation payoff $\widehat{\alpha}$. Therefore his choices in equilibrium must maximize the

[14] The model focuses on the reputational concerns of CEOs, rather than the private benefits CEOs stand to gain upon certain events transpiring, or upon profit sharing in the event the investment goes well. Examples of private benefits include large payouts or "golden parachutes" that become payable upon a change of control transaction, and stock in the private company that accrues to management in the event of a management buyout (MBO) of a publicly traded company. Examples of profit sharing include grants of company stock and stock options. While in some cases, private benefits and profit sharing will be large enough to overcome concerns about reputation, it is legitimate to think that in general, efforts to influence standing with the board and with the corporate community at large will be a guiding motivation of many CEOs.

expectation of the continuation payoff over all possible realizations of investment outcomes. This structure is analogous to a game in which there are two periods and play ends after the second period. The CEO makes choices to maximize the expected posterior belief of the board.

Figure 1a

Board mandates and separated reporting

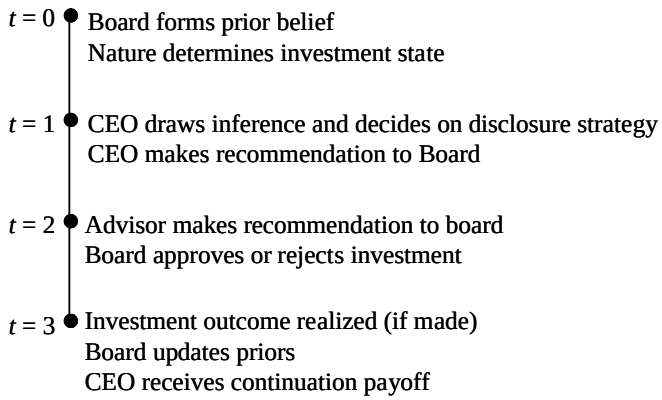

Figure 1b

Board mandates and pooled reporting

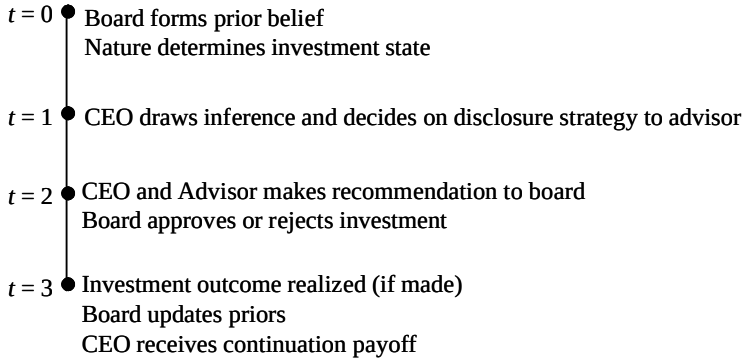

Figure 1c

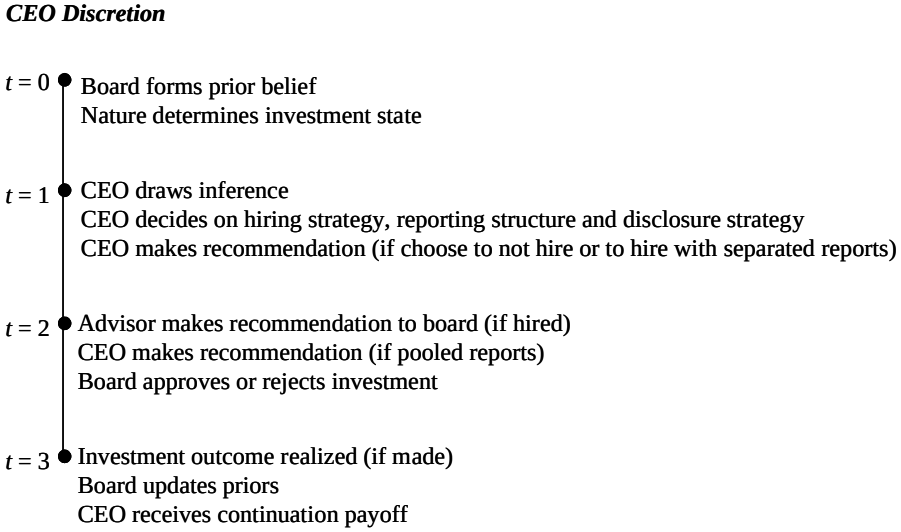

3 Equilibrium Results

Throughout the analysis, I utilize the equilibrium solution concept of Perfect Bayesian Nash Equilibrium. To obtain uniqueness, I impose the intuitive criterion (see Cho and Kreps (1987)). The intuitive criterion eliminates equilibria that can only be sustained by imposing unreasonable beliefs about out-of-equilibrium behavior.[15]

Before analyzing the solution to this game, it is instructive to provide a roadmap outlining the results that follow. Table 2 summarizes the equilibrium results. The equilibrium analysis proceeds by first examining CEO incentives to reveal the full information set to the advisor if hired. After establishing the disclosure strategy (which depends on the probability of detection), I discuss the equilibrium outcomes in the following four scenarios:

i) Board mandates use of advisor when probability of detection is "high"

ii) CEO discretion over use of advisor when probability of detection is "high"

iii) Board mandates use of advisor when probability of detection is "low"

iv) CEO discretion over use of advisor when probability of detection is "low"

[15] In signaling games, the intuitive criterion requires that if the information set following a certain action a_j by a given type t_i is off the equilibrium path, and that action is equilibrium-dominated for that type, then the receiver's belief $\Pr(t_i \mid a_j)$ should place zero probability on the player being type t_i. In a perfect Bayesian equilibrium in a signaling game, the action a_j is *equilibrium-dominated* for type t_i if t_i's equilibrium payoff is greater than t_i's highest possible payoff from playing a_j. See Cho and Kreps (1987).

Table 2: Summary of Equilibrium Results

	Probability of detection "high"	**Probability of detection "low"**
Board Mandate	With separated reporting, the advisor's recommendation assists board in both approval and monitoring duties. With pooled reporting, the advisor's recommendation assists board in approval but not monitoring duty.	With separated reporting, the advisor's recommendation does not assist board in either approval or monitoring duties. With pooled reporting, the advisor's recommendation assists board in approval but not monitoring duty.
CEO Discretion	CEO always hires advisor and chooses separated reporting. Advisor's recommendation assists board in both approval and monitoring duties.	CEO only hires advisor after bad recommendation. Advisor's recommendation does not assist board in either approval or monitoring duties.

The analysis is divided into mandated use and discretionary use of advisors to reflect heterogeneity of use in practice. In general, I define "mandate" to encompass any instance where the board seeks a second opinion. For example, audit committees on boards of publicly traded companies are required to hire external auditors to provide audits of corporate financial disclosures[16]; investment banks provide fairness opinions[17]; and attorneys perform due diligence for firm disclosures.[18] On the other hand, "discretion" refers to management control over the use of outside experts. For example, management consultants are hired by, and work directly with management to form a recommendation on a given strategy; investment banks bring investment opportunities to firms, or will

[16] SOX requires external auditors to report directly to the audit committee (section 301). The Securities Act of 1933 and Securities and Exchange Act of 1934 requires issuers of publicly traded securities to audit their financial statements.

[17] *Smith vs. Van Gorkum* (488 A. 2d. 875-878 (1985)) is the seminal case establishing that second opinions, known as "fairness opinions", serve as a liability shield for boards from shareholder lawsuits. In *Van Gorkum*, the Delaware Supreme Court stated that while the obtainment of an independent fairness opinion in a change of control transaction is not required, it is evidence that directors have met the enhanced fiduciary burden that arises when transactions are not conducted at arms length. Cases subsequent to *Van Gorkum* have confirmed that fairness opinions issued by investment banks, consulting firms, accounting firms or other independent advisor, are sufficient to effectively remove legal risk for gross negligence for board directors.

[18] The Securities Act of 1933 imposed strict liability on any issuer of securities that makes a material misstatement or omission in the disclosure document. As a defense, officers and directors need to show that they conducted a "reasonable investigation" verifying the accuracy of claims made in the issuing document. The obligation to conduct a reasonable investigation has essentially been delegated to the company's corporate attorneys. See Coffee (2006).

work with management to assess potential acquisitions or merger targets; and attorneys and accountants work with in-house counsel and finance department on legal and accounting issues respectively.

3.1 CEO Disclosure Strategy

To solve for the equilibrium of this game, we use backward induction starting with the disclosure decision, since this is the CEO's final strategic node.[19] The CEO wishes to maximize his expected continuation payoff $E[\widehat{\alpha}]$. Therefore, the disclosure strategy depends on how the board updates its belief about ability type. If recommendations do not match, the board's posterior is always $\widehat{\alpha}(\widehat{s} \neq s_a) = 0$. Since the high type draws an identical inference to the advisor, he will always prefer to reveal the full information set. This implies $\varphi_H\left(U \mid A, a^h, r^s\right) = \varphi_H\left(U \mid A, a^h, r^s\right) = 0$. The low type, who does not know what the advisor will recommend, must decide whether to reveal his signal, risking the chance that the advisor will disagree with him, or to strategically withhold elements, risking detection by the advisor. The following two lemmas outline the disclosure strategies for the low type under pooled and separated reporting structures respectively.

Lemma 1 *(**Disclosure strategy with pooled reports**) When CEO and advisor reports are pooled, low types will reveal the true ambiguous signal to the advisor using a disclosure strategy* $\varphi_L\left(U \mid A, a^h, r^p\right) = 0$.

Proof. Since reports are pooled, the low type sees the advisor's recommendation prior to making his own. Because the advisor always agree with the high type (because they have perfectly correlated inferences), the low type will also always make the same recommendation as the advisor. Since recommendations never disagree when reports are pooled, the board's posterior belief that the CEO has high ability is $\widehat{\alpha} = \alpha$. With partial disclosure there is a risk that the CEO will be detected censoring elements of the signal by the advisor, and so the expected continuation payoff is $\rho\alpha \leq \alpha$. Therefore, the low type maximizes his expected payoff when $\varphi_L\left(U \mid A, a^h, r^p\right) = 0$. ■

Because there is full disclosure, the advisor's recommendation will be reliable. By contrast, the disclosure strategy when reports are separated depends directly on the probability of detection. To simplify notation, and to account for the fact that strategies may be different depending on the recommendation, I substitute $\varphi_L\left(\widehat{s}^{\theta}\right)$ for $\varphi_L\left(U \mid A, a^h, r^p\right)$. Letting $x \equiv p\theta + (1-p)(1-\theta)$, the probability the high type makes a good recommendation, and recalling that $\sigma_L\left(\widehat{s}^{G}\right)$denotes the probability that the low type makes a good recommendation, with separated reports, the board's posterior following good recommendations

[19] Even though the game continues after this point, board approval decisions and updating of priors that take place after information disclosure are purely mechanical.

by both the CEO and advisor, is:

$$\begin{aligned} & \widehat{\alpha}\left(\widehat{s}^{G}, s_{a}^{G} \mid a^{h}, r^{s}, \varphi\right) \\ = & \frac{\alpha}{\alpha+(1-\alpha)\left(\varphi_{L}\left(\widehat{s}^{G}\right)+\left(1-\varphi_{L}\left(\widehat{s}^{G}\right)\right) x \sigma_{L}\left(\widehat{s}^{G}\right)\right)} \end{aligned} \tag{1}$$

The low type uses a disclosure strategy $\varphi_{L}\left(\widehat{s}^{G}\right)$ that maximizes his expected continuation payoff conditional upon a good recommendation:

$$\begin{aligned} & E_{L}\left[\widehat{\alpha} \mid \widehat{s}^{G}, a^{h}, r^{s}\right] \\ = & \left(\varphi_{L}\left(\widehat{s}^{G}\right)(1-\rho)+\left(1-\varphi_{L}\left(\widehat{s}^{G}\right)\right) x\right) \widehat{\alpha}\left(\widehat{s}^{G}, s_{a}^{G} \mid a^{h}, r^{s}, \varphi\right) \end{aligned} \tag{2}$$

Following a bad recommendation by both the CEO and the advisor, the board's posterior is:

$$\begin{aligned} & \widehat{\alpha}\left(\widehat{s}^{B}, s_{a}^{B} \mid a^{h}, r^{s}, \varphi\right) \\ = & \frac{\alpha}{\alpha+(1-\alpha)\left(\varphi_{L}\left(\widehat{s}^{B}\right)+\left(1-\varphi_{L}\left(\widehat{s}^{B}\right)\right)(1-x)\left(1-\sigma_{L}\left(\widehat{s}^{G}\right)\right)\right)} \end{aligned} \tag{3}$$

The expected continuation payoff for the low type, conditional upon a bad recommendation is therefore:

$$\begin{aligned} & E_{L}\left[\widehat{\alpha} \mid \widehat{s}^{G}, a^{h}, r^{s}\right] \\ = & \left(\varphi_{L}\left(\widehat{s}^{B}\right)(1-\rho)+\left(1-\varphi_{L}\left(\widehat{s}^{B}\right)\right)(1-x)\right) \widehat{\alpha}\left(\widehat{s}^{B}, s_{a}^{B} \mid a^{h}, r^{s}, \varphi\right) \end{aligned} \tag{4}$$

Defining $\rho^{t\left(\widehat{s}^{\theta}\right)}$ as the *cutoff detection threshold*, we have the following lemma.

Lemma 2 *(**Disclosure strategy with separated reports**): When $\rho<\rho^{t\left(\widehat{s}^{\theta}\right)}$, the low type always discloses an unambiguous signal ($\varphi_{L}\left(\widehat{s}^{\theta}\right)=1$). When $\rho>\rho^{t\left(\widehat{s}^{\theta}\right)}$, the low type always discloses the true ambiguous signal ($\varphi_{L}\left(\widehat{s}^{\theta}\right)=0$). The cutoff detection threshold $\rho^{t\left(\widehat{s}^{\theta}\right)}$ is strictly increasing in α. Further, $\rho^{t\left(\widehat{s}^{G}\right)}$ is strictly decreasing in x, and $\rho^{t\left(\widehat{s}^{B}\right)}$ is strictly increasing in x.*

The proof is contained in the Appendix. Lemma 2 establishes that when reports are separated, the disclosure strategy depends directly on the likelihood that third party advisors detect that the signal they receive from management is incomplete. When the probability of detection is "low" (that is, when $\rho<\rho^{t\left(\widehat{s}^{\theta}\right)}$), the CEO presents an unambiguous signal, favoring the recommendation he has already made. The risk that the advisor does not certify the CEO's recommendation with full disclosure outweighs the risk of being caught with partial disclosure. When the probability of detection is "high" ($\rho>\rho^{t\left(\widehat{s}^{\theta}\right)}$), the CEO presents an ambiguous signal.

The cutoff threshold detection level $\rho^{t\left(\widehat{s}^{\theta}\right)}$ depends on the overall state of the economy θ. When the investment outcome is likely to be good, the advisor

is more likely to make a good recommendation. The CEO then has less need to withhold information, since he has greater certainty over how the advisor will report. The same logic applies when the outcome is likely to be bad. Disclosing the true signal poses less of a risk, since recommendations are likely to match. Cutoff detection thresholds increase when the investment outcome is less certain since the low type will be unsure of the advisor's recommendation and therefore more willing to lie.

The cutoff detection threshold also increases with the board's prior belief about CEO ability α. Low types held in relatively high regard by the board, have more to lose if the advisor fails to provide certification. Therefore, to protect their reputation, they are more likely to withhold information to ensure certification, even as the risk of detection increases. The escalating riskiness of the behavior engaged in by executives at Enron and WorldCom exemplifies this point. Even as the likelihood of detection grew, executives continued to manipulate and withhold information from the board and other third parties. The costs of exposure were so great that seemingly reckless behavior became rational.[20]

The hope of legislative efforts, such as SOX, is to increase the incentives of management to fully disclose relevant information to third parties. Below I show that whether SOX is effective, depends directly on whether these efforts are successful. The results above demonstrate that such efforts are important since low types will withhold information from certifying parties when it is not likely they will be detected doing so. The cutoff detection thresholds are:

$$\rho^{t\left(\widehat{s}^G\right)} = \frac{\alpha + x\sigma_L\left(\widehat{s}^G\right) - \alpha x\sigma_L\left(\widehat{s}^G\right) - x}{\alpha + x\sigma_L\left(\widehat{s}^G\right) - \alpha x\sigma_L\left(\widehat{s}^G\right)}$$

$$\rho^{t\left(\widehat{s}^B\right)} = \frac{\alpha\left(x + \sigma_L\left(\widehat{s}^G\right) + x\sigma_L\left(\widehat{s}^G\right)\right) - (1-x)\,\sigma_L\left(\widehat{s}^G\right)}{\alpha\left(x + \sigma_L\left(\widehat{s}^G\right) + x\sigma_L\left(\widehat{s}^G\right)\right)}$$

Notice that because the detection thresholds are strictly less than one, detection does not have to be perfect to ensure management cooperation. Therefore full auditing of every management claim is unnecessary. If increasing the probability of detection is costly, boards can invest just to the point at which management would rather disclose the true signal.

We are now in a position to continue with the equilibrium analysis. Given that disclosure will either be full or partial, we can likewise split up the analysis. Since the focus of the paper is on the impact of regulation, we can further divide into the discussion into examination of mandated versus discretionary hiring decisions.

[20] Kofman and Lawarrée (1993) reach a similar conclusion noting that collusion with internal auditors (with side payments) becomes more attractive for managers the larger their punishment is for low effort. Also see Chevalier and Ellison (1997) for empirical evidence that mutual fund managers with poor performances choose riskier portfolios than those with good performances.

3.2 High probability of detection if CEO censors private information

3.2.1 Board mandates use of advisors

In circumstances where the probability of detection is high enough ($\rho > \rho^{t(\widehat{s}^{\theta})}$), low types will always fully disclose their signal regardless of the reporting structure. When reports are pooled, even though the certification is reliable, no learning about CEO type is possible. High types will always report their true inference. Low types who see the advisor's recommendation will simply make the same recommendation to the board.

With separated reporting, the low type does not have an opportunity to see the advisor's recommendation before making his own, and therefore will not know how the advisor will recommend. He does however know the likelihood the advisor draws good and bad inferences, and how the board updates following certification. A Perfect Bayesian Nash Equilibrium is defined below.

Definition 3 ***(Perfect Bayesian Nash Equilibrium with mandates and full information)****: i) Knowing his own ability type and after drawing an inference from his private signal, the CEO utilizes a recommendation strategy $\sigma(\widehat{s} \mid s)$ that maximizes his expected continuation payoff $E[\widehat{\alpha}]$; ii) after observing the CEO's strategy and using the advisor's recommendation, the board updates its prior belief about CEO ability using Bayes law; iii) the board approves investments to maximize expected firm profits; iv) board beliefs are consistent with CEO strategy and CEO strategy is optimal given board beliefs.*

Because inferences are imperfectly correlated with the true investment state, it is possible for advisors to be wrong when acting honestly and relying on the true signal. Even so, obtaining an informed recommendation increases the likelihood that the board makes decisions that maximize profits, and thus it will always follow the advisor's recommendation.[21] For the CEO, all that matters is whether or not the advisor certifies the CEO's recommendation, rather than the actual investment outcome if realized. If the advisor fails to certify the CEO's recommendation, the board infers that the CEO has low ability. Substituting $\varphi_L\left(s^{\theta}\right) = 0$ into the updating equations above (see equations (1) and (3)), given that the high type always reports his true inference (because he knows that the advisor is honest, and will always agree with him), the board's posteriors following a good recommendation and a bad recommendation, both certified by the advisor respectively are:

$$\widehat{\alpha}\left(\widehat{s}^{G}, s_a^{G}\right) = \frac{\alpha}{\alpha + (1-\alpha)\, x \sigma_L\left(\widehat{s}^{G}\right)} \tag{5}$$

$$\widehat{\alpha}\left(\widehat{s}^{B}, s_a^{B}\right) = \frac{\alpha}{\alpha + (1-\alpha)\,(1-x)\left(1-\sigma_L\left(\widehat{s}^{G}\right)\right)} \tag{6}$$

[21] Courts have cited the fact that a board has obtained a second opinion from an independent advisor about a potential investment, as evidence that it has fulfilled its procedural fiduciary duty to shareholders. See *Smith v Van Gorkum*.

With mandates in place, the only thing the CEO must decide is the recommendation to make to the board. The low type will set $\sigma_L\left(\widehat{s}^G\right)$ to maximize his expected continuation payoff:

$$E_L\left[\widehat{\alpha}\right] = \sigma_L\left(\widehat{s}^G\right) x\widehat{\alpha}\left(\widehat{s}^G, s_a^G\right) + \left(1 - \sigma_L\left(\widehat{s}^G\right)\right)(1-x)\,\widehat{\alpha}\left(\widehat{s}^B, s_a^B\right)$$

Proposition 4 details the equilibrium reporting strategies. The proofs for this and all other propositions are contained in the Appendix below.

Proposition 4 ***(Mandates with full information and separated reports)*** *When advice is mandated, reports are separated and* $\rho > \rho^{t\left(\widehat{s}^\theta\right)}$*, there exists a unique Perfect Bayesian Nash Equilibrium where i) the high type truthfully reports his inference; ii) the low type reports a good recommendation with probability* $\sigma_L^*\left(\widehat{s}^G\right)$*, where* $\sigma_L^*\left(\widehat{s}^G\right) = 1$ *for* $\theta \geq \overline{\theta}$*,* $\sigma_L^*\left(\widehat{s}^G\right) = 0$ *for* $\theta \leq \underline{\theta}$*,* $x < \sigma_L^*\left(\widehat{s}^G\right) < 1$ *for* $\frac{1}{2} < \theta < \overline{\theta}$*,* $0 < \sigma_L^*\left(\widehat{s}^G\right) < x$ *for* $\underline{\theta} < \theta < \frac{1}{2}$*, and* $\sigma_L^*\left(\widehat{s}^G\right) = x = \frac{1}{2}$ *for* $\theta = \frac{1}{2}$*; and iii) the board approves the investment when the advisor makes a good recommendation and* $\Pr\left(G \mid s_a^G\right) - c > 0$*.*

When signals are accurately and fully disclosed, certification of management decisions can be of great assistance to boards for several reasons. Although costly[22] and perhaps of limited direct informational value (as happens when the CEO has high ability), certification increases the likelihood of good corporate investments, enables boards to learn faster about CEO ability, and to not unduly punish high types for unforeseen investment outcomes. High types recommend truthfully, knowing that the advisor will certify his recommendation. Low types employ a strategy that maximizes the chances that he will recommend in the same way as the advisor, utilizing a pure strategy of always reporting good (bad) if the unconditional likelihood the investment is good (bad) is high (low) enough. Otherwise he uses a mixed strategy, overemphasizing the recommendation the advisor is more likely to make. This result provides some justification for mandates, and is no doubt, the outcome policy makers have in mind when proposing regulations that increase board independence, and require external certification of management recommendations.

3.2.2 CEO discretion over use of advisors

Without certification, the board uses its observation of the recommendation and the investment outcome (if known) to update its priors. The posterior following a good recommendation and good outcome is,

$$\widehat{\alpha}\left(\widehat{s}^G, G\right) = \frac{\alpha p}{\alpha p + (1-\alpha)\,\sigma_L\left(\widehat{s}^G\right)} \tag{8}$$

[22] Implicitly assumed throughout the paper is that the advisor's fee z does not exceed the expected return net of investment costs. If advisory fees are too costly, then boards would prefer never to obtain certification. Note that if the CEO has discretion, even when fees exceed the net return, he will still obtain certification because it is the company who pays the fee not the CEO. The CEO will use company resources inefficiently to increase his own private payoff. Requiring the CEO to pay the advisor directly may result in a separating equilibrium because hiring only makes economic sense for high types.

The posterior following a good recommendation and bad outcome is,

$$\widehat{\alpha}\left(\widehat{s}^{G}, B\right) = \frac{\alpha\left(1-p\right)}{\alpha\left(1-p\right) + \left(1-\alpha\right)\sigma_{L}\left(\widehat{s}^{G}\right)} \tag{9}$$

After a bad recommendation (where no outcome is observed), the posterior is,

$$\widehat{\alpha}\left(\widehat{s}^{B}\right) = \frac{\alpha\left(1-x\right)}{\alpha\left(1-x\right) + \left(1-\alpha\right)\left(1-\sigma_{L}\left(\widehat{s}^{G}\right)\right)} \tag{10}$$

CEOs choose hiring, reporting and recommendation strategies to maximize expected continuation payoffs. The payoff to different hiring and reporting strategies for the high type depends directly on the low type's recommendation strategy $\sigma_{L}\left(\widehat{s}^{G}\right)$. Given this strategy we can calculate equilibrium payoffs for the high type under different equilibria (whose existence depends directly on board beliefs about out-of-equilibrium behavior). Utilizing the intuitive criterion to eliminate equilibria that require unreasonable beliefs about out-of-equilibrium behavior, I demonstrate that there is a unique equilibrium where both CEO types always hire and choose a separated reporting structure.

Proposition 5 *(**CEO discretion with full information**): There exists a unique Perfect Bayesian Nash equilibrium in which both types always hire an advisor, choose a separated reporting structure, and give recommendations in accordance with Proposition 4 above. The board believes the CEO has low ability if he chooses not to hire an advisor, or if he chooses a pooled reporting structure. The board approves investments when the advisor makes a good recommendation, and* $\Pr\left(G \mid s_{a}^{G}\right) - c > 0$.

Giving the CEO discretion over when to obtain a second opinion, results in exactly the same equilibrium as described in Proposition 4. Because certification is a signal to the board of CEO ability, by seeking certification, high types can actually improve their expected reputations relative to the case when no advisor is available. The uniqueness of this equilibrium arises from the fact that expected payoffs following certification strictly dominate any other equilibrium payoffs for high types, but leave low types worse off than were they not to hire or pool reports. This result is robust for all values of the board's prior α. The fact that certification provides additional information to the board about CEO ability, is especially important since with no certification, given the low type's strategy, high types only improve their reputation when they make a good recommendation and the investment outcome is good. Otherwise the posteriors decline relative to the prior.

In an equilibrium where the CEO makes a good recommendation without using an advisor, I show in the Appendix below that the low type makes a good recommendation more often than the high type is wrong – that is, $\sigma_{L}\left(\widehat{s}^{G}\right) > 1-p$. This implies that $\widehat{\alpha}\left(\widehat{s}^{G}, B\right) < \alpha$; not surprisingly, reputation declines after a

positive recommendation and a bad outcome. On the other hand, given that investment outcomes are unobservable when no investment is made, without advisors low types are more inclined to reject investments than high types. Therefore, $1 - \sigma_L\left(\widehat{s}^G\right) = \sigma_L\left(\widehat{s}^B\right) < x$ implying that $\widehat{\alpha}\left(\widehat{s}^B\right) < \alpha$: following an uncertified bad recommendation, reputation also declines. Certification therefore protects high types from adverse inferences. Low types also seek certification, even though they are strictly worse off, since the board believes upon seeing out-of-equilibrium behavior, that the CEO has low ability.

To support this result, we have evidence that there was strong demand for external auditing in the United States in the 1920's, prior to the institution of auditing mandates. By 1933, 85% of all publicly traded corporations already sought outside audits as a signal to investors of disclosure quality.[23] More recently, there is evidence that such signaling occurred with the new SOX requirement that CEOs certify the company's audited statements. Executives at several public companies (including among others, Fannie Mae, Gannett Co, Marriott, Corning, Delphi Corp) certified their financial disclosures before the signing deadline as a way to demonstrate their confidence in their statements. J.T. Battenberg, CEO of Delphi Corp and one of the first executives to certify, stated, "Some CEOs said I put them on the defensive." Similarly Corning Inc stated that it wanted to certify when it was not yet required to do so, because it would underscore its differences with other firms in the telecommunications sector.[24]

The results in Propositions 4 and 5 are dependent on the following assumptions: (1) managers are mainly concerned about their reputations; (2) managers are aware that they are being monitored by their superiors; and (3) advisors are assured access to all relevant information. If managers stand to gain some private benefit from a particular decision by the board, obtaining certification may not be part of an equilibrium strategy under a discretionary regime. For example, if the private inference for an investment opportunity is bad, but the CEO receives a very large payout in the event the investment is made – big enough to outweigh concerns about potential reputational damage – then high types will not choose to seek certification on their own, preferring to recommend in favor of the investment. Similarly, if low types do not believe they will be duly punished, they would rather not obtain external certification because it increases their risk of exposure. Mandates may be useful devices in both of these cases *so long as the information disclosed by management is accurate*. If the information is unreliable, mandates can be counterproductive elevating the status of costly second opinions that are ultimately uninformative. The remain-

[23] Colonel A.H. Carter, President of New York Society of Certified Public Accountants in testimony to the Senate Committee on Banking and Currency, 73rd Cong. 56-62 (1933). In these same hearings Senator Gore pointed out that pervasive auditing failed to prevent the 1929 Crash. See Previts and Merino (1998).

[24] Downey Grimsley, "Signing the Bottom Line; Top Executives at 16 Companies Certify Their Books", *Washington Post*, August 1, 2002.

der of the paper investigates outcomes when management disclosures may be incomplete.

3.3 Low probability of detection if CEO censors information

3.3.1 Board mandates use of advisors

When the probability of detection is small enough, low types will always present partial unambiguous signals to advisors. This result, contained in Lemma 2, leads naturally to Proposition 6.

Proposition 6 *(**Mandates with separated reporting and partial disclosure**): When advice is mandated, reports are separated and $\rho < \rho^{t\left(\widehat{s}^{\theta}\right)}$, i) high types will disclose the true ambiguous signal and truthfully recommend in accordance with their inferences; ii) low types will disclose an unambiguous signal, recommending in accordance with the signal presented. The board approves the investment when the advisor makes a good recommendation, and $\alpha \Pr\left(G \mid \widehat{s}^{G}\right) + (1-\alpha)(1-\rho)\theta - c > 0$.*[25]

With separated reports and low detection probabilities, the advisor will either certify the CEO's recommendation – often incorrectly – or will discover on occasion that the CEO has misrepresented his signal. The incentive to censor elements in the signal arises from the separated reporting structure. Since the low type is required to recommend prior to the advisor, he will give the advisor information that leads her recommendation to match his. Since we assume that $\rho < \rho^{t\left(\widehat{s}^{\theta}\right)}$, we have $\varphi_L\left(\widehat{s}^{G}\right) = 1$ by Lemma 2. Substituting into equation (1), the board's posterior belief is therefore identical to its prior[26]:

$$\alpha\left(\widehat{s}^{G}, s_a^{G} \mid a^h, r^s, \rho < \rho^{t\left(\widehat{s}^{G}\right)}\right) = \alpha$$

The posterior with separated reporting is therefore equal to the posterior with pooled reporting, where there is no disincentive to full information disclosure (by Lemma 1). Low types present the true signal, matching their recommendation to that of the advisor. The advisor's recommendation provides valuable information about the investment, but not about the CEO. Proposition 7 (which relies on Lemma 1 for disclosure strategies) summarizes CEO strategies when detection probability is low, and reports are pooled.

Proposition 7 *(**Mandates with pooled reporting and low detection probabilities**): When advice is mandated, reports are pooled and $\rho < \rho^{t\left(\widehat{s}^{G}\right)}$, i) high types truthfully recommend in accordance with their inferences, and ii) low types, after seeing the advisor's recommendation, will make identical recommendations. The board approves the investment when the advisor makes a good recommendation and $\Pr\left(G \mid \widehat{s}^{G}\right) - c > 0$.*

[25] Note that with certification, since the CEO's payoff is not related to investment outcome, the low type is indifferent between making a good recommendation and a bad recommendation.

[26] Note that the same point applies for certification after a bad recommendation.

Proposition 7 demonstrates that boards may not be able to carry out both monitoring and approval functions when CEOs have career concerns.[27] Information transmission by self-interested CEOs has been investigated by Adams and Ferreira (2006) and Song and Thakor (2006). Adams and Ferreira come to the conclusion that firms with friendly boards (who do not monitor CEOs too intensely) may have better investment outcomes than firms with independent boards. The CEO's trade-off between presenting good information, and being held accountable for poor recommendations, means that he is not inclined to present information to the board which could later be held against him. The results here, which take into account the role of third parties in assisting boards, lead one to a similar conclusion. Boards that monitor too closely may be given inferior advice.

3.3.2 CEO discretion over use of advisors

With discretion over hiring decisions, high types will take into account the incentives of low types to present unambiguous signals to the advisor. Because the board is unable to update its prior when the advisor always agrees with the CEO, certification loses its signaling power. However, high types still hire sometimes to avoid adverse inferences about ability type following a bad recommendation. The equilibrium is described in Proposition 8.

Proposition 8 *(**CEO discretion with partial information**): When outside advice is discretionary, and* $\rho < \rho^{t\left(\widehat{s}^G\right)}$*, the following strategies and beliefs constitute a unique equilibrium: i) conditional on a good inference, the high type never hires an advisor and always makes a good recommendation; ii) conditional on a bad inference, the high type always hires an advisor, is indifferent between reporting structures and makes a bad recommendation; iii) the low type makes a good recommendation and does not hire with probability* $\sigma_L^d\left(\widehat{s}^G\right) \in (1-p, x)$*; (iv) the low type makes a bad recommendation and hires an advisor with probability* $1-\sigma_L^d\left(\widehat{s}^G\right)$*; (v) the board believes that the CEO has low ability if he hires an advisor, and the advisor makes a good recommendation; (vi) the board approves the investment when the CEO makes a good recommendation and* $\alpha \Pr\left(G \mid \widehat{s}^G\right) + (1-\alpha)\,\theta - c > 0$*.*

Conditional on a good inference, high types will maximize their expected continuation payoffs by letting the board evaluate their ability based on investment performance alone. The expected continuation payoff following a good

[27] Note however, that pooled reports may not remedy CEO incentives to partially disclose information to the advisor if the CEO stands to gain some private benefit from a particular investment recommendation.

recommendation is:

$$\begin{aligned} E_H\left[\widehat{\alpha} \mid \widehat{s}^G, a^n\right] &= \Pr\left(G \mid \widehat{s}^G, H\right)\widehat{\alpha}\left(\widehat{s}^G, G\right) + \Pr\left(B \mid \widehat{s}^G, H\right)\widehat{\alpha}\left(\widehat{s}^G, B\right) \\ &= \frac{p\theta}{p\theta + (1-p)(1-\theta)}\frac{\alpha p}{\alpha p + (1-\alpha)\sigma_L^d\left(\widehat{s}^G\right)} + \\ &\quad \frac{(1-p)(1-\theta)}{p\theta + (1-p)(1-\theta)}\frac{\alpha(1-p)}{\alpha(1-p) + (1-\alpha)\sigma_L^d\left(\widehat{s}^G\right)} \end{aligned}$$

Substituting for $\sigma_L^d\left(\widehat{s}^G\right)$, we can show that $E_H\left[\widehat{\alpha} \mid \widehat{s}^G, a^n\right] > \alpha$[28]: the high type is better off not hiring an advisor after drawing a good inference. However, seeking certification after drawing a bad inference is still desirable for high types for the same reasons as described above. Since low types reject more often than high types in equilibrium, certification is a way to protect reputation.[29]

The low type utilizes a mixed recommendation and hiring strategy. Always rejecting investments (and hiring) cannot be an equilibrium strategy. Upon seeing a good recommendation, the board would think the CEO had high ability regardless of the outcome. The low type will therefore deviate and make a good recommendation. The same rationale applies for why always recommending in favor of the investment (and not hiring), is not an equilibrium strategy either.

Collusion between agents and supervisors is not necessary to generate results where certifications are used to deceive principals. The results from Propositions 6 and 8 demonstrate why a mandate can be a very blunt mechanism that may ultimately fail. With discretion, actors can take their own and others' incentives into account when deciding on their strategy. Actions themselves can be valuable signaling mechanisms. Removing the capacity to act so as to signal can exacerbate the very behavior hoped to be deterred by regulation.

4 Conclusion

Management has incentives to distort the private information it discloses when it knows that its performance is being actively evaluated. Oversight therefore

[28] This is formally proved in the proof to Proposition 5 contained in the Appendix below.

[29] The model assumes that the board is aware of the investment opportunity. In practice, the board will very often rely upon the CEO to bring investment opportunities to its attention. Since low types reject investments more often than do high types, there is an adverse reputational inference upon a bad recommendation regardless of type. Therefore, both high types and low types would prefer to simply not inform the board that the investment exists in the first place. If the board is unaware of the investment, it will not expect a recommendation, and no updating occurs. In the case where boards are unaware of investment opportunities, with low detection probabilities, neither type seeks certification, and boards only ever receive good recommendations. In this case, there is no reason for high types to hire advisors when the probability of detection is low. However, with high detection probabilities, high types will still seek certification for bad recommendations since the expected continuation payoff $\frac{\alpha}{\alpha+(1-\alpha)(1-x)\sigma_L^d\left(\widehat{s}^B\right)}$ is strictly greater than α.

creates the very problem it was designed to overcome. Profit sharing may ameliorate and private benefits may exacerbate the incentives to disclose full information. Nevertheless, so long as management cares enough about its reputation, it will disclose information to advisors strategically to improve its standing with the board, not always in accordance with the best interests of shareholders.

For this reason, institutional rules that co-exist with the obligation to monitor (or that increase the independence of boards) assume great significance. In this paper, I have examined the circumstances in which external advisors can assist boards in carrying out their fiduciary duties to shareholders to maximize firm profits, and to monitor management. The ability to monitor depends squarely on the advisor's ability to gain unfettered access to internal information. Rules that increase the incentives for management to fully disclose all that it knows directly facilitate monitoring. On the other hand, rules that require certification without taking into account these incentives will be redundant at best, and detrimental at worst.

Given that management's reputation depends so heavily on the information it discloses, and that outside certification can a valuable monitoring and approval device, the question naturally arises as to why firms don't institute systems on their own to ensure adequate disclosures and outside certification. Separation between ownership and control combined with dispersed ownership, means that entrenched managers will choose contracts (or charters and by-laws) that benefit them directly, possibly at the expense of shareholders who have little ability to exert control (Berle and Means (1932), Cary (1974), Coffee (1988)). It is thought that mandates are necessary to overcome problems involved with ensuring that companies enact efficient contracts. Instituting requirements that increase the ability and likelihood that boards monitor runs counter to the self-interest of some CEOs, meaning that changes to by-laws which would implement these measures are unlikely. However, so long as CEOs retain control over internal corporate information, the value of mandates overriding these contractual arrangements may not be of great value.

The costs of ensuring the quality of private information disclosures (indirectly incorporated in the model into the advisory fee z) lie at the heart of the criticisms of SOX. Indeed there is empirical evidence to suggest that companies at the margin, for whom the cost of improving information quality exceeds the benefit from better investment decisions and certifications, have decided to opt out of the public markets altogether (Engel, Hayes and Wang (2004)). By escaping provisions regarding the quality of internal control systems, these companies also escape mandates requiring increased monitoring. This may not be such a bad thing, because when information is too "soft" or too costly to verify, monitoring can generate the wrong incentives. In this paper, I show that mandates increasing the intensity of oversight are only desirable when coupled with measures that increase the likelihood that management will fully disclose all relevant information.

5 Appendix

Proof. *(Lemma 2)* If the low type decides to make a good recommendation, he sets his disclosure strategy $\varphi_L\left(\widehat{s}^G\right)$ to maximize his expected continuation payoff. Taking a first derivative of equation (1) with respect to $\varphi_L\left(\widehat{s}^G\right)$, we have:

$$\frac{\partial E_L\left[\widehat{\alpha} \mid \widehat{s}^G, a^h, r^s\right]}{\partial \varphi_L\left(\widehat{s}^G\right)} = \alpha - x - \alpha\rho - x\alpha\rho + x\varphi_L\left(\widehat{s}^G\right)(1 - \alpha + \alpha\rho) \tag{11}$$

The sign of equation (11) depends on the probability of detection ρ. We can see that the marginal return from lying is increasing for values of that satisfy:

$$\rho < \frac{\alpha + x\sigma_L\left(\widehat{s}^G\right) - \alpha x\sigma_L\left(\widehat{s}^G\right) - x}{\alpha + x\sigma_L\left(\widehat{s}^G\right) - \alpha x\sigma_L\left(\widehat{s}^G\right)} \equiv \rho^{t\left(\widehat{s}^G\right)} \tag{12}$$

The low type maximizes (1) by making $\varphi_L\left(\widehat{s}^G\right)$ as large as possible. Therefore, he sets $\varphi_L\left(\widehat{s}^G\right) = 1$, and only discloses an unambiguously good signal. When $\rho > \rho^{t\left(\widehat{s}^G\right)}$, equation (1) is decreasing, and the expected continuation payoff is maximized be making $\varphi_L\left(\widehat{s}^G\right)$ as small as possible. Therefore, the low type always discloses the true ambiguous signal, letting $\varphi_L\left(\widehat{s}^G\right) = 0$.

We can similarly establish the detection threshold when the low type wishes to report a bad recommendation. The marginal return from lying is increasing for values of such that:

$$\rho < \frac{\alpha\left(x + \sigma_L\left(\widehat{s}^G\right) + x\sigma_L\left(\widehat{s}^G\right)\right) - (1 - x)\,\sigma_L\left(\widehat{s}^G\right)}{\alpha\left(x + \sigma_L\left(\widehat{s}^G\right) + x\sigma_L\left(\widehat{s}^G\right)\right)} = \rho^{t\left(\widehat{s}^B\right)} \tag{13}$$

and decreasing for $\rho > \rho^{t\left(\widehat{s}^B\right)}$. To maximize the expected continuation payoff conditional on making a bad recommendation, the low type will set $\varphi_L\left(\widehat{s}^G\right) = 1$ for $\rho < \rho^{t\left(\widehat{s}^B\right)}$, and $\varphi_L\left(\widehat{s}^G\right) = 0$ for $\rho > \rho^{t\left(\widehat{s}^B\right)}$. ■

Proof. *(Proposition* 4) The proof proceeds by assuming that the equilibrium exists and then demonstrating that CEO strategies are optimal given board beliefs. Since advisors are mandated, the only relevant belief is the updated posterior. The high type will always recommend in accordance with his inference, since his inference is perfectly correlated with the advisor. The low type does not know what the advisor will report (remembering that the signal is ambiguous), but does know the likelihood the advisor draws and reports good and bad inferences. Making a good recommendation always ($\sigma_L\left(\widehat{s}^G\right) = 1$) is optimal so long as the expected continuation payoff from doing so, exceeds the payoff from deviating and reporting a bad signal. Substituting into equations (5) and (6), recognizing that $\sigma_L\left(\widehat{s}^B\right) = 1 - \sigma_L\left(\widehat{s}^G\right)$, and that the advisor reports a good inference with probability x and a bad inference with probability $1 - x$, the condition necessary to sustain such a pure strategy is

$$x\frac{\alpha}{\alpha + (1 - \alpha)\,x} \geq 1 - x \tag{14}$$

The LHS of equation (14) is the expected continuation payoff from always making a good recommendation while the RHS is the expected payoff from deviating by sometimes making a bad recommendation. Define $\overline{\theta} < 1$ as the point that (14) binds with equality. For values of $\theta > \overline{\theta}$, the low type optimizes his expected payoff by setting $\sigma_L^*\left(\widehat{s}^G\right) = 1$. On the other hand, reporting a bad recommendation always, $\sigma_L^*\left(\widehat{s}^G\right) = 0$, is only optimal if the expected payoff from doing so exceeds the expected payoff from deviating. The necessary condition to sustain such a pure strategy is

$$(1-x)\frac{\alpha}{\alpha + (1-\alpha)(1-x)} \geq x \tag{15}$$

As before, the LHS of equation (15) is the payoff from always making a bad recommendation, while the RHS is the payoff from deviating. Define $\underline{\theta} > 0$ as the point that equation (15) binds with equality. For values of $\theta < \underline{\theta}$, the low type optimizes his expected payoff by always making a bad recommendation, $\sigma_L^*\left(\widehat{s}^G\right) = 0$. For $\underline{\theta} < \theta < \overline{\theta}$, it will not be possible to sustain a pure strategy equilibrium since the low type is better off deviating. Therefore $0 < \sigma_L^*\left(\widehat{s}^G\right) < 1$. Note that equation (5) is decreasing with $\sigma_L^*\left(\widehat{s}^G\right)$, while equation (**??**) is increasing with $\sigma_L^*\left(\widehat{s}^G\right)$. This means that the low type maximizes his expected payoff when satisfies the condition:

$$x\frac{\alpha}{\alpha + (1-\alpha)\, x\sigma_L^*\left(\widehat{s}^G\right)} = (1-x)\frac{\alpha}{\alpha + (1-\alpha)(1-x)\left(1-\sigma_L^*\left(\widehat{s}^G\right)\right)} \tag{16}$$

Solving explicitly for $\sigma_L^*\left(\widehat{s}^G\right)$, we can show that for $\frac{1}{2} < \theta < \overline{\theta}$, $\sigma_L^*\left(\widehat{s}^G\right) > x$; and for $\underline{\theta} < \theta < \frac{1}{2}$, $\sigma_L^*\left(\widehat{s}^G\right) < x$. Only for $\theta = \frac{1}{2}$, does the low type make good and bad recommendations with exactly the same probability as the high type: $\sigma_L^*\left(\widehat{s}^G\right) = x = \frac{1}{2}$. ■

Before proving Proposition 5 directly, we must consider other possible equilibrium outcomes. Lemmas 9 – 12 provide details of strategies and beliefs in equilibria in which advisors are not always hired.

Lemma 9 *(**Pooled Reporting**) There exists an equilibrium in which both types always hire an advisor and choose pooled reports The high type recommends in accordance with his inference, and the low type always agrees with the advisor. If the CEO decides not to hire or chooses separated reports, the board believes he is the low type. The board approves investments when the advisor makes a good recommendation and* $\Pr\left(G \mid s_a^G\right) - c > 0$.

Proof. Given board beliefs, both types will always hire and choose pooled reports. The high type will report truthfully since his inference is perfectly correlated with the advisor. The low type will always recommend in the same way as the advisor because he sees the advisor's recommendation prior to making his own. Given that the board cannot compare the recommendations, it's posterior (regardless of the recommendation is) $\widehat{\alpha}\left(a^h, r^s\right) = \alpha$. Since the advisor has

access to full information, the board will maximize period profits by following the advisor's recommendation so long as the net expected benefit is positive. ■

Lemma 10 *(**No Advisors**) There exists an equilibrium in which i) neither type ever hires an advisor; ii) the high type always recommends in accordance with his inference; iii) low type uses a mixed recommendation strategy, making a good recommendation with probability* $\sigma_L^n\left(\widehat{s}^G\right) \in (1-p, x)$. *If the CEO chooses to hire an advisor, the board believes he has low ability.*

Proof. Suppose this equilibrium exists. Since the low type draws an uninformative inference, the probability that he will be correct following a good recommendation is just θ. His expected continuation payoff from making a good recommendation is therefore:

$$E_L\left[\widehat{\alpha} \mid \widehat{s}^G\right] = \theta\widehat{\alpha}\left(\widehat{s}^G, G\right) + (1-\theta)\,\widehat{\alpha}\left(\widehat{s}^G, B\right) \tag{17}$$

Following a bad recommendation is expected payoff is,

$$E_L\left[\widehat{\alpha} \mid \widehat{s}^B\right] = \widehat{\alpha}\left(\widehat{s}^B\right) \tag{18}$$

In equilibrium, the CEO will set $\sigma_L\left(\widehat{s}^G\right)$ to maximize his total expected payoff:

$$E_L\left[\widehat{\alpha}\right] = \sigma_L\left(\widehat{s}^G\right) E_L\left[\widehat{\alpha} \mid \widehat{s}^G\right] + \left(1 - \sigma_L\left(\widehat{s}^G\right)\right) E_L\left[\widehat{\alpha} \mid \widehat{s}^B\right] \tag{19}$$

Since $E_L\left[\widehat{\alpha} \mid \widehat{s}^G\right]$ is decreasing in $\sigma_L\left(\widehat{s}^G\right)$, and $E_L\left[\widehat{\alpha} \mid \widehat{s}^B\right]$ is increasing in $\sigma_L\left(\widehat{s}^G\right)$, will be maximized when satisfies the following condition:

$$E_L\left[\widehat{\alpha} \mid \widehat{s}^G\right] = E_L\left[\widehat{\alpha} \mid \widehat{s}^B\right] \tag{20}$$

If $\sigma_L\left(\widehat{s}^G\right) = 0$, then $E_L\left[\widehat{\alpha} \mid \widehat{s}^G\right] > E_L\left[\widehat{\alpha} \mid \widehat{s}^B\right]$. This means that always making a bad recommendation cannot be part of an equilibrium strategy since deviation is more profitable. Likewise, if $\sigma_L\left(\widehat{s}^G\right) = 1$, $E_L\left[\widehat{\alpha} \mid \widehat{s}^G\right] < E_L\left[\widehat{\alpha} \mid \widehat{s}^B\right]$: again deviating is optimal with a strategy which requires the CEO to always make a good recommendation. Therefore, the low type will always use a mixed recommendation strategy. We can check that when $\sigma_L\left(\widehat{s}^G\right) = x$ (the probability the high type makes a good recommendation), $E_L\left[\widehat{\alpha} \mid \widehat{s}^G\right] < E_L\left[\widehat{\alpha} \mid \widehat{s}^B\right]$. The low type will want to place more weight on bad recommendations relative to the high type. Likewise, if $\sigma_L\left(\widehat{s}^G\right) = x$, we have $E_L\left[\widehat{\alpha} \mid \widehat{s}^G\right] > E_L\left[\widehat{\alpha} \mid \widehat{s}^B\right]$. The low type will want to deviate and make a good recommendation. By continuity, there exists some $\sigma_L^n\left(\widehat{s}^G\right) \in (1-p, x)$ which satisfies equation (20).

The expected continuation payoff for the high type following a good recommendation is:

$$E_H\left[\widehat{\alpha} \mid \widehat{s}^G\right] = \Pr\left(G \mid \widehat{s}^G, H\right)\widehat{\alpha}\left(\widehat{s}^G, G\right) + \Pr\left(B \mid \widehat{s}^G, H\right)\widehat{\alpha}\left(\widehat{s}^G, B\right) \tag{21}$$

Following a bad recommendation, the expected continuation payoff is just $\widehat{\alpha}\left(\widehat{s}^B\right)$, the same as the low type. For truth-telling to be an optimal strategy for the high type, the following incentive constraints must be satisfied:

$$\Pr\left(G \mid \widehat{s}^G, H\right)\widehat{\alpha}\left(\widehat{s}^G, G\right) + \Pr\left(B \mid \widehat{s}^G, H\right)\widehat{\alpha}\left(\widehat{s}^G, B\right) > \widehat{\alpha}\left(\widehat{s}^B\right) \tag{22}$$

$$\widehat{\alpha}\left(\widehat{s}^B\right) > \Pr\left(G \mid \widehat{s}^B, H\right)\widehat{\alpha}\left(\widehat{s}^G, G\right) + \Pr\left(B \mid \widehat{s}^B, H\right)\widehat{\alpha}\left(\widehat{s}^G, B\right) \tag{23}$$

Note that $\sigma_L^n\left(\widehat{s}^G\right) < x$ implies that $\sigma_L^n\left(\widehat{s}^G\right) < p$, and hence $\widehat{\alpha}\left(\widehat{s}^G, G\right) > \alpha > \widehat{\alpha}\left(\widehat{s}^G, B\right)$. Since by definition, $\widehat{\alpha}\left(\widehat{s}^B\right)$ is a convex combination of $\widehat{\alpha}\left(\widehat{s}^G, G\right)$ and $\widehat{\alpha}\left(\widehat{s}^G, B\right)$, it must be that $\widehat{\alpha}\left(\widehat{s}^G, G\right) > \widehat{\alpha}\left(\widehat{s}^B\right) > \widehat{\alpha}\left(\widehat{s}^G, B\right)$. Since by assumption, $p > 1-p$, $\Pr\left(G \mid \widehat{s}^G, H\right) > \theta$ and $\Pr\left(B \mid \widehat{s}^G, H\right) < 1-\theta$. Therefore equation (22) is always satisfied. Since $\Pr\left(G \mid \widehat{s}^B, H\right) < \theta$ and $\Pr\left(B \mid \widehat{s}^B, H\right) > 1-\theta$, equation (23) is also satisfied. Hence, truth-telling is optimal for the high ability CEO. ■

Lemma 11 *(**Hire after good inference only**) There exists an equilibrium in which i) the high type always recommends in accordance with his inference, and chooses to hire an advisor only after a good inference; ii) the low type utilizes a pure strategy of only making bad recommendations when $\theta \leq \theta^-$, and a mixed strategy making a good recommendation with probability $\sigma_L^{nb}\left(\widehat{s}^G\right) < x$ when $\theta > \theta^-$. If the CEO chooses to hire an advisor, and the advisor makes a bad recommendation, the board believes the CEO has low ability.*

Proof. Suppose this equilibrium exists. Using the same logic as outlined above, the low type chooses to maximize his expected payoff which is,

$$E_L\left[\widehat{\alpha}\right] = \sigma_L\left(\widehat{s}^G\right)x\widehat{\alpha}\left(\widehat{s}^G, s_a^G\right) + \left(1 - \sigma_L\left(\widehat{s}^G\right)\right)\widehat{\alpha}\left(\widehat{s}^B\right) \tag{24}$$

It is optimal for the low type to use a pure recommendation strategy ($\sigma_L\left(\widehat{s}^G\right) = 0$) when,

$$x < \frac{\alpha\left(1-x\right)}{\alpha\left(1-x\right)+1-\alpha} \tag{25}$$

Let be such that equation (25) binds with equality. For $\theta \leq \theta^-$, the equilibrium recommendation strategy is $\sigma_L^{nb}\left(\widehat{s}^G\right) = 0$. For $\theta > \theta^-$, $\sigma_L^{nb}\left(\widehat{s}^G\right)$ must satisfy the condition:

$$x\widehat{\alpha}\left(\widehat{s}^G, s_a^G\right) = \widehat{\alpha}\left(\widehat{s}^B\right) \tag{26}$$

Solving explicitly for the equilibrium recommendation strategy , we can show that $\sigma_L^{nb}\left(\widehat{s}^G\right) < x$. Truth-telling is always optimal for the high type: making a bad recommendation after drawing a good inference results in no certification, and reporting a bad signal is dominated by obtaining certification since $\widehat{\alpha}\left(\widehat{s}^G, s_a^G\right) > \widehat{\alpha}\left(\widehat{s}^B\right)$. ■

Lemma 12 *(**Hire after bad inference only**) There exists an equilibrium in which i) the high type always recommends in accordance with his inference, and chooses to hire an advisor only after a bad inference; ii) the low type utilizes a pure strategy of only making good recommendations when $\theta \geq \theta^+$, and a mixed strategy making a good recommendation with probability $\sigma_L^{ng}\left(\widehat{s}^G\right) > x$ when $\theta < \theta^+$. If the CEO chooses to hire an advisor, and the advisor makes a good recommendation, the board believes the CEO has low ability.*

Proof. Suppose this equilibrium exists. Again, using the same logic as above, the low type chooses $\sigma_L\left(\widehat{s}^G\right)$ to maximize his expected payoff which is,

$$E_L\left[\widehat{\alpha}\right] = \sigma_L\left(\widehat{s}^G\right) E_L\left[\widehat{\alpha} \mid \widehat{s}^G\right] + \left(1 - \sigma_L\left(\widehat{s}^G\right)\right)(1-x)\,\widehat{\alpha}\left(\widehat{s}^B, s_a^B\right) \tag{27}$$

It is optimal for the CEO to use a pure strategy of always recommending in favor of the investment ($\sigma_L\left(\widehat{s}^G\right) = 1$) so long as,

$$1 - x < \theta \frac{\alpha p}{\alpha p + 1 - \alpha} + (1-\theta)\frac{\alpha(1-p)}{\alpha(1-p) + 1 - \alpha} \tag{28}$$

Let θ^+ be such that equation (28) binds with equality. For $\theta \geq \theta^+$, the equilibrium recommendation strategy is $\sigma_L^{ng}\left(\widehat{s}^G\right) = 1$. For $\theta < \theta^+$, $\sigma_L^{ng}\left(\widehat{s}^G\right)$ must satisfy the condition:

$$E_L\left[\widehat{\alpha} \mid \widehat{s}^G\right] = (1-x)\,\widehat{\alpha}\left(\widehat{s}^B, s_a^B\right) \tag{29}$$

Solving explicitly for the equilibrium recommendation strategy , it is possible to show that $\sigma_L^{ng}\left(\widehat{s}^G\right) < x$. Truth telling is always optimal since $\widehat{\alpha}\left(\widehat{s}^B, s_a^B\right) > \Pr\left(G \mid s^B\right)\widehat{\alpha}\left(\widehat{s}^G, G\right) + \Pr\left(B \mid s^B\right)\widehat{\alpha}\left(\widehat{s}^G, B\right)$, and $E_H\left[\widehat{\alpha} \mid \widehat{s}^G, a^n\right] > 0$. ■

We are now in a position to prove that an equilibrium in which the CEO always chooses to hire an advisor exists, and that it is unique by virtue of equilibrium dominant payoffs.

Proof. *(Proposition 5)* Assume that this equilibrium exists. By Proposition 4, if he hires an advisor to whom he gives the full information set and chooses a separated reporting strategy, the low type maximizes his expected payoff by using the recommendation strategy $\sigma_L^*\left(\widehat{s}^G\right)$. Thus, the expected payoff for the low type is,

$$E_L\left[\widehat{\alpha} \mid a^h, r^s\right] = \sigma_L^*\left(\widehat{s}^G\right) x\widehat{\alpha}\left(\widehat{s}^G, s_a^G\right) + \left(1 - \sigma_L^*\left(\widehat{s}^G\right)\right)(1-x)\,\widehat{\alpha}\left(\widehat{s}^B, s_a^B\right) \tag{30}$$

If the low type deviates from the proposed equilibrium strategies by either not hiring an advisor, or by choosing pooled reports, given board beliefs, his expected payoff is,

$$E_L\left[\widehat{\alpha} \mid a^n\right] = E_L\left[\widehat{\alpha} \mid a^h, r^p\right] = 0 \tag{31}$$

Since $E_L\left[\widehat{\alpha} \mid a^h, r^s\right] > 0$, deviating yields a lower expected continuation payoff. Therefore the proposed strategies and beliefs constitute an equilibrium.

For uniqueness, we need to show that the expected payoff for the high type equilibrium dominates all other expected payoffs from other possible equilibria. The payoff from always hiring an advisor with a separated reporting structure is:

$$E_H\left[\widehat{\alpha} \mid a^h, r^s\right] = x\widehat{\alpha}\left(\widehat{s}^G, s_a^G\right) + (1-x)\,\widehat{\alpha}\left(\widehat{s}^B, s_a^B\right) \tag{32}$$

The equilibrium payoff following pooled reports is clearly dominated by equation (32). Because $x\sigma_L^*\left(\widehat{s}^G\right) < 1$, we have $\widehat{\alpha}\left(\widehat{s}^G, s_a^G\right) = \frac{\alpha}{\alpha + (1-\alpha)x\sigma_L^*\left(\widehat{s}^G\right)} > \alpha$. Because $(1-x)\left(1 - \sigma_L^*\left(\widehat{s}^G\right)\right) < 1$, we have $\widehat{\alpha}\left(\widehat{s}^B, s_a^B\right) = \frac{\alpha}{\alpha + (1-\alpha)(1-x)\left(1-\sigma_L^*\left(\widehat{s}^G\right)\right)} > \alpha$.

Now we can eliminate equilibria where no advisor is used. The total expected payoffs from always using advisors, never using advisors (Lemma 10), using advisors only after a good inference (Lemma 11) and only using an advisor after a bad inference is (Lemma 12) respectively are:

$$E_H\left[\widehat{\alpha} \mid a^n\right] = p\theta\widehat{\alpha}\left(\widehat{s}^G, G\right) + (1-p)(1-\theta)\,\widehat{\alpha}\left(\widehat{s}^G, B\right) + (1-x)\,\widehat{\alpha}\left(\widehat{s}^B\right) \quad (33)$$

$$E_H\left[\widehat{\alpha} \mid a^{h|\widehat{s}^G}\right] = x\widehat{\alpha}\left(\widehat{s}^G, G\right) + (1-x)\,\widehat{\alpha}\left(\widehat{s}^B\right) \quad (34)$$

$$E_H\left[\widehat{\alpha} \mid a^{h|\widehat{s}^B}\right] = p\theta\widehat{\alpha}\left(\widehat{s}^G, G\right) + (1-p)(1-\theta)\,\widehat{\alpha}\left(\widehat{s}^G, B\right) + (1-x)\,\widehat{\alpha}\left(\widehat{s}^B\right) \quad (35)$$

Comparing, the expected payoff from always hiring strictly dominates all other expected continuation payoffs. Therefore, beliefs necessary to sustain other equilibria (that only low ability CEOs deviate and hire) are unreasonable. The equilibrium where advisors are always hired is therefore unique. ■

Proof. *(Proposition 8)* Suppose this equilibrium exists. Given that the high type reports truthfully after a good inference, the expected payoff for the low type conditional on making a good recommendation is:

$$E_L\left[\widehat{\alpha} \mid \widehat{s}^G\right] = \theta\widehat{\alpha}\left(\widehat{s}^G, G\right) + (1-\theta)\,\widehat{\alpha}\left(\widehat{s}^G, B\right) \quad (36)$$

The expected payoff for the low type following a bad recommendation (given that the low type will present an unambiguously bad signal to the advisor),

$$E_L\left[\widehat{\alpha} \mid \widehat{s}^B\right] = \alpha(1-\rho) \quad (37)$$

The low type chooses $\sigma_L\left(\widehat{s}^G\right)$ to maximize his total expected payoff,

$$E_L[\widehat{\alpha}] = \sigma_L\left(\widehat{s}^G\right) E_L\left[\widehat{\alpha} \mid \widehat{s}^G\right] + \left(1 - \sigma_L\left(\widehat{s}^G\right)\right) E_L\left[\widehat{\alpha} \mid \widehat{s}^B\right] \quad (38)$$

Equation (36) is decreasing in $\sigma_L\left(\widehat{s}^G\right)$ while equation (37) is constant in $\sigma_L\left(\widehat{s}^G\right)$. Setting , we can solve for the equilibrium recommendation strategy that maximizes the low type's payoff. By the same methodology as used above, we can show that $\sigma_L^d\left(\widehat{s}^G\right) \in (1-p, x)$. ■

References

[1] Adams, Renée and Daniel Ferreira (2006), "A Theory of Friendly Boards", *Journal of Finance*, 62, 217-250.

[2] Aghion Philippe and Jean Tirole (1997), "Formal and Real Authority in Organizations", *Journal of Political Economy*, 105, 1-29.

[3] Almazan, Andres and Javier Suarez (2003), "Entrenchment and severance pay in optimal governance", *Journal of Finance*, 58, 519-547.

[4] Baron, David P. and David Besanko (1984), "Regulation, Asymmetric Information, and Auditing", *RAND Journal of Economics*, 15, 447-470.

[5] Bebchuk, Lucian, Alma Cohen and Allen Ferrell (2005), "What matters in corporate governance?", Discussion Paper 491, John M. Olin Center for Law Economics and Business, Harvard Law School.

[6] Bebchuk, Lucian and Marcel Kahan (1989), "Fairness Opinions: How fair are they and what can be done about it?", *Duke Law Journal*, 1989, 27-53.

[7] Berle Adolf, and Gardiner Means, T*he Modern Corporation and Private Property*, 1932.

[8] Bhagat, Sanjai and Bernard Black (1999), "The Uncertain Relationship Between Board Composition and Firm Performance", *Business Lawyer*, 54, 921-963.

[9] Bowers, Helen (2002) "Fairness Opinions and the Business Judgment Rule: An empirical investigation of target firms' use of fairness opinions", *Northwestern University Law Review*, 96.

[10] Bowers, Helen and William Latham III, " Information asymmetry, litigation risk, uncertainty and the demand for fairness opinions: Evidence from US mergers and acquisitions, 1980-2002", Working Paper, (April 13, 2006).

[11] Cary, William (1974), "Federalism and Corporate Law: Reflections Upon Delaware", 83 *Yale Law Journal,* 663.

[12] Chevalier, Judith and Glenn Ellison (1997), "Risk Taking by Mutual Funds as a Response to Incentives", *Journal of Political Economy*, 105.

[13] Chharochharia, Vidhi and Yaniv Grinstein, "The Transformation of US Corporate Boards: 1997-2003", Working Paper (May 2004).

[14] Cho, In-Koo and David Kreps (1987), "Signaling games and stable equilibria", *Quarterly Journal of Economics*, 102, 179-221.

[15] Coates, John. C (2007), "The Goals and Promises of Sarbanes-Oxley", *Journal of Economic Perspectives*, 91.

[16] Coffee, John (1988), "No Exit? Opting Out, the Contractual Theory of the Corporation, and the Special Case of Remedies", 53 *Brookings Law Review* 919.

[17] Coffee, John (2006), *Gatekeepers: The Professions and Corporate Governance*, Oxford University Press, New York.

[18] Cotter, James, Anil Shivdasani and Marc Zenner (1997), "Do independent directors enhance target shareholder wealth during tender offers?", *Journal of Financial Economics*, 43, 195-218.

[19] Cox, James D. (2003), "Managing and Monitoring Conflicts of Interest: Empowering the Outside Directors with Independent Counsel", *Villanova Law Review*, 48, 1077-1095.

[20] Demski, Joel (2003), "Corporate Conflicts of Interest", *Journal of Economic Perspectives*, 17, 51-72.

[21] Ely, Jeffrey and Juuso Valimaki (2003), "Bad Repuation", *Quarterly Journal of Economics*, 118, 785-814.

[22] Fairfax, Lisa (2002), "Form over Substance? Officer Certification and the Promise of Enhanced Personal Accountability under the Sarbanes-Oxley Act", 55 *Rutgers Law Review* 1.

[23] Fama, Eugene and Michael Jensen (1983), "Separation of Ownership and Control", *Journal of Law and Economics*, 26, 301-325.

[24] Fama, Eugene and Michael Jensen (1983), "Agency Problems and Residual Claims", *Journal of Law and Economics*, 26, 327-349.

[25] Frankel, Michael (2005), *Mergers and Acquisitions Basics*, John Wiley and Sons, New Jersey.

[26] Fudenberg, Drew and David K. Levine (1989), "Reputation and Equilibrium Selection in Games with a Patient Player", *Econometrica*, 57, 759-778.

[27] Hazard, Geoffrey and Edward B. Rock, "A new player in the Boardroom: The emergence of the independent directors' counsel", University of Pennsylvania Law School, Institute for Law and Economics (May 2004).

[28] Hermalin, Benjamin and Michael Weisbach (1988), "The Determinants of Board Composition", *RAND Journal of Economics*, 589-606.

[29] Hermalin, Benjamin and Michael Weisbach (1998), "Endogenously Chosen Boards of Directors and Their Monitoring of the CEO", *American Economic Review*, 88, 96-118.

[30] Holmstrőm, Bengt (1999), "Managerial Incentive Problems: A Dynamic Perspective", Review of Economic Studies, 66, 169-182, originally in *Essays in Economics and Management in Honor of Lars Wahlbeck*, Helsinki: Swedish School of Economics (1982).

[31] Holmstrőm, Bengt and Steven Kaplan (2001), "Corporate Governance and Merger Activity in the United States: Making Sense of the 1980s and 1990s", *Journal of Economic Perspectives*, 15, 121-144.

[32] Holmstrőm, Bengt and Joan Ricart i Costa (1986), "Managerial Incentives and Capital Management", *Quarterly Journal of Economics*, 101, 835-860.

[33] Khalil, Fahad (1997), "Auditing without Commitment", *RAND Journal of Economics*, 28, 629-640.

[34] Kisgen, Darren, Jun Qian and Weihong Song, "Are fairness opinions fair? The case of mergers and acquistions", Working Paper (January 28, 2007).

[35] Kofman, Fred and Jaques Lawarrée (1993), "Collusion in Hierarchical Agency", *Econometrica*, 61, 629-656.

[36] Kraakman, Reinier H. (1986), "Anatomy of a Third-Party Enforcement Strategy", *Journal of Law, Economic and Organization*, 2, 53-104.

[37] Kreps, David M. and Robert Wilson (1982), " Reputation and Imperfect Information", *Journal of Economic Theory*, 27, 253-279.

[38] Levitt, Steven and Christopher Snyder (1997), "Is no news bad news? Information transmission and the role of "early warning" in the principal-agent model", *RAND Journal of Economics*, 28, 641-661.

[39] Makhija, Anil and Rajesh Narayanan, "Fairness Opinions in Mergers and Acquisitions", Working Paper (April 11, 2007).

[40] Morris, Stephen (2001), "Political Incorrectness", *Journal of Political Economy*, 109, 231-265.

[41] Ottoviani, Marco and Peter Sorenson (2000), "Herd Behavior and Investment: Comment", *American Economic Review*, 90, 695-704.

[42] Ribstein, Larry (2002), "Market vs Regulatory Responses to Corporate Fraud: A Critique of the Sarbanes-Oxley Act of 2002", 28 *Journal of Corporate Law*.

[43] Romano, Roberta (2005), "The Sarbanes-Oxley Act and the Making of Quack Corporate Governance", 114 *Yale Law Journal* 1521.

[44] Scharfstein, David and Jeremy Stein (1990), "Herd Behavior and Investment", *American Economic Review*, 80, 465-479.

[45] Song, Fenghua and Anjan Thakor (2006), "Information Control, Career Concerns and Corporate Governance", *Journal of Finance*, 61, 1845-1896.

[46] Trueman, Brett, (1994), "Analyst Forecasts and Herding Behavior", *The Review of Financial Studies*, 7, 97-124.

[47] Veasey, E. Norman (2003), "Corporate Governance and Ethics in the Post-Enron Worldcom Environment", *Wake Forest Law Review*, 38, 839-853.